arXiv:1511.07536v1 [cs.LO] 24 Nov 2015

A Symbolic Logic with Concrete Bounds for Cryptographic Protocols

Anupam Datta[A], Joseph Y. Halpern[B], John C. Mitchell[C], Arnab Roy[D] and Shayak Sen[A]

[A]Computer Science Department, Carnegie Mellon University
[B]Computer Science Department, Cornell University
[C]Department of Computer Science, Stanford University
[D]Fujitsu Laboratories of America

Abstract

We present a formal logic for quantitative reasoning about security properties of network protocols. The system allows us to derive concrete security bounds that can be used to choose key lengths and other security parameters. We provide axioms for reasoning about digital signatures and random nonces, with security properties based on the concrete security of signature schemes and pseudorandom number generators (PRG). The formal logic supports first-order reasoning and reasoning about protocol invariants, taking concrete security bounds into account. Proofs constructed in our logic also provide conventional asymptotic security guarantees because of the way that concrete bounds accumulate in proofs. As an illustrative example, we use the formal logic to prove an authentication property with concrete bounds of a signature-based challenge-response protocol.

1 Introduction

Large and complex cryptographic protocols form the backbone of internet security today. At the scale of the internet, protocols use keys to communicate with millions of agents over extended periods of time, with increasingly powerful adversaries attempting to break them. In this scenario, evaluating security in a quantitative sense with respect to key lengths and adversary runtimes is critical. Concrete security is a practice-oriented approach to cryptography that requires that proofs of a security property be accompanied by a concrete probability with which the property is false.

We present Quantitative Protocol Composition Logic (QPCL), a formal system that supports concrete security proofs of probabilistic safety properties of cryptographic protocols. A typical assertion in QPCL is a concrete security

statement of the form $B^{\epsilon}(\varphi)$, which we call a *probabilistic belief*. Here φ is a temporal safety property of the execution of a protocol in the presence of probabilistic polynomial time adversaries, and ϵ is a probability that is a function of η, the security parameter governing the length of keys and nonces, and t, the concrete runtime of the adversary. This assertion is true if φ holds with probability at least $1 - \epsilon(\eta, t)$. Typically, ϵ also depends on functions that represent the probability of an adversary violating the security of the specific cryptographic primitives used in a protocol.

QPCL's proof system provides a separation among axioms and rules for cryptographic primitives, probabilistic beliefs, and programs. QPCL axioms about cryptographic primitives such as signature schemes and pseudorandom-number generators introduce a probabilistic error bound. These axioms have formal soundness proofs that show that whenever an axiom is false on a trace, an attack on the corresponding primitive can be constructed. Therefore, these probabilistic error bounds are closely related to the security of the underlying primitive. These soundness proofs are a central contribution of this paper. Axioms about cryptographic primitives introduce concrete bounds that are combined in QPCL using proof rules for reasoning about probabilistic beliefs, adapting an approach introduced by Halpern [1], which in turn is based on the *ϵ-semantics* of Goldszmidt, Morris, and Pearl [2]. Traces generated by protocol programs serve as an instance of the more general semantic models of Halpern, allowing us to adapt the sound and complete proof system from [1] into our logic. Finally, in the style of Protocol Composition Logic, QPCL's predecessor for symbolic [3] and asymptotic reasoning [4], QPCL includes a Hoare-style program modality, and an operator for temporal ordering of actions that allows reasoning about non-monotonic safety properties of program traces.

Reasoning about concrete bounds in QPCL necessitates some important differences from its symbolic and asymptotic predecessors. Firstly, to achieve precise concrete bounds, the placement of quantifiers must be carefully considered. For example, $\forall x.B^{\epsilon}(\varphi)$ expresses a probabilistic bound on φ that applies to all choices of x, whereas $B^{\epsilon}(\forall x.\varphi)$ bounds the probability of $\forall x.\varphi$; when x ranges over a large domain, a more precise bound can be derived for the first formula. Secondly, we restrict the proof system so that only safety properties are provable in QPCL. This restriction allows us to follow a useful proof strategy, where we first reason about the probability of a post-condition of a program being false, and then transfer a bound on that probability to the end of the execution of the protocol. For safety properties the probability of a formula being false at some point in an execution is bounded by the probability of a formula being false at the end of that execution. Section 5.4 discusses the importance of the restriction to safety properties.

QPCL assertions are interpreted over traces generated by a simple probabilistic, concurrent programming language executing in the presence of an adversary. The semantics of the programming language is independent of any particular cryptographic primitive, due to a separation between actions and their implementations. Our formal treatment of this separation is simpler than a full module system such as the one used in [5], and yet exposes the features

(runtimes and trace probabilities) needed to define the semantics of probabilistic trace properties. This separation also allows the adversary to be represented by any program that runs in polynomial time, therefore allowing the soundness proofs of our logic to be valid in a computational model.

We illustrate the use of the QPCL reasoning principles by proving a "matching-conversations" style authentication property [6] for a simple initiator-responder protocol, which states that after the protocol is completed, it is known that the precise order in which messages are sent and received by the participants of the protocol was not altered by an adversary. Section 6 contains an outline of the QPCL proof.

Prior work on formal proof techniques for concrete cryptography [7, 8, 9] has largely focused on relational techniques where security is proved via game-based reductions, similar to traditional cryptographic proofs. Proofs of a number of cryptographic schemes have been mechanized using these techniques. However, we conjecture that reasoning about the security of cryptographic protocols involves relatively simpler reasoning about trace properties, and that relational reasoning, while important, does not need to be exposed to the user of a proof system. In prior work [10], indistinguishability properties such as secrecy have been proven via stronger trace properties such as key secrecy. Therefore, relational reasoning in QPCL is pushed down to the level of the soundness proofs of axioms, which involve a one-time cost to the designers of the proof system; the users of QPCL need to reason only about trace properties of protocols.

The rest of this paper is organized as follows. Section 2 is a brief introduction to concrete security. Section 3 is an overview of the programming language and the proof system. The programming language for protocols and its operational semantics is described in Section 4. The syntax, semantics, and proof system for QPCL are described in Section 5. Section 6 has an outline of the QPCL proof of the initiator-responder protocol considered in the paper. Section 7 contains a comparison of QPCL with closely related work.

2 Concrete Security Primer

Concrete security proofs provide bounds for the probability with which an attacker can subvert some security notion for a primitive. These bounds usually are a function of

1. the security parameter η, which determines the length of keys and random numbers;[1]

2. the runtime bound of the adversary, which is often expressed as a polynomial $tb(\eta)$ of the security parameter; and

3. the number of attempts available to the adversary to break the primitive. As we describe below, the exact definition of the third parameter is

[1] For ease of presentation, in this paper, we assume that every cryptographic primitive is governed by the same security parameter η.

dependent on the definition of security for that primitive.

In this paper, we study the concrete security of protocols that involve digital signatures and pseudorandom-number generators; we briefly describe them here.

A *digital signature scheme* is designed to ensure message unforgeability in a public key setting. An agent is assigned a private signing key which has an associated public verification key. A signature scheme $\Pi = (Gen, Sign, Verify)$ is specified by a key generation algorithm, Gen, a signing algorithm, $Sign$, and a verification algorithm $Verify$. Gen generates pairs of public and private keys, $Sign$ produces signatures over messages, and $Verify$ checks that a signature was generated using the signing key. A characterization of security for a signature scheme Π is unforgeability under chosen message attacks [11], shortened to $\mathsf{uf\text{-}cma}$, which requires that the probability of the adversary creating a new message/signature pair, given that he can see q signatures produced by Π over messages of his choice, is bound by some negligible function $\epsilon_{\Pi}^{\mathsf{uf\text{-}cma}}(\eta, tb, q)$.

A *pseudorandom-number generator* P is an algorithm that produces numbers that are distinguishable from numbers drawn from a uniformly random distribution by a computationally bounded adversary with only negligible probability. This probability is called the advantage of the adversary against a scheme P, and we denote it by $\epsilon_{P}^{\mathsf{prg}}(\eta, tb, q)$, where q is the number of samples that the adversary sees before deciding which distribution the samples are drawn from. Psuedorandom number generators are often used in protocols to generate secrets that are hard to guess by the adversary.

3 Overview of the Formal System

QPCL is a probabilistic program logic with concrete security bounds for temporal trace properties of cryptographic protocols. A concrete security bound on an assertion in our logic expresses the probability with which a computationally constrained adversary can violate the assertion in terms of the length of keys, the runtime of the adversary and the number of concurrent sessions of the protocol.

3.1 Programming Model

The models on which the logic is interpreted are traces generated by a simple protocol programming language executing in the presence of an adversary. The requirements of the logic neccessitate some features of the operational semantics of the protocol programming language, as the operational semantics defines how traces are generated from protocol programs interacting in the presence of an adversary. We now motivate some of these features.

- *Concrete Runtime.* Reasoning about concrete runtime requires accounting for the runtime of every computation in an execution. This runtime includes that of the adversary as well as protocol programs. The oper-

ational semantics is defined by a transition system where transitions are labeled by these runtime costs.

- *Concrete Probabilities.* Computing probabilistic bounds requires considering probabilities of sets of traces. The transition system is probabilistic and the probabilistic branches induce an execution tree, where each path in the tree is an execution trace. We use the tree to compute the relevant probabilities.

- *Computational completeness.* For proofs to be valid in a computational model, we cannot restrict the adversary to a symbolic model. However, modeling the adversary explicitly in a Turing-complete language makes the language very complicated. For proofs of security, which are generally proofs via reduction to known hard problems, it is possible to ignore the representation of the adversary. Therefore, we do not explicitly restrict the adversary's program to a particular language, but reason only about its input-output behavior.

- *Concurrency and Adversarial Scheduling.* Programs execute concurrently; where the scheduler that switches between programs is controlled by the adversary. The cost of the computation required by the adversary to make scheduling decision is accounted for in the runtime of a trace.

3.2 Logic and Proof System

QPCL allows declarative specifications of temporal trace properties, and the proof system supports reasoning in a manner very similar to first-order reasoning. Essentially, axioms related to cryptography have error bounds associated with them, and these error bounds are propagated throughout the proof to derive a bound on the final assertion. To aid reasoning about cryptographic protocols, QPCL supports the following features:

- *Temporal Trace Properties.* Protocols, such as those involving authentication, often require properties that specify the precise order in which certain events occur. To allow reasoning about properties such as message ordering, QPCL supports temporal operators. Traditionally, such properties are specified using games [12], which imply that the security proofs typically involve the harder task of proving equivalence of games.

- *Concrete Security Bounds.* Concrete security bounds on assertions in QPCL depend on the protocol being reasoned about in a number of ways. First, protocols specify the implementations of cryptographic primitives to be used; these define the concrete bounds on the axioms related to cryptography. Second, concrete bounds on the cryptographic primitives usually also depend on the number of queries made to the primitive, and to bound this number on any particular trace, we use the number of times the cryptographic primitive is called in a protocol.

(bitstrings)	b	::=	$\hat{A}$	principal name
			n	nonce
			c	string constant
			s	signature
			ι	session identifier
(thread identifier)	I	::=	$\langle \hat{A}, \iota \rangle$	thread ι of agent $\hat{A}$
(terms)	t	::=	x	variable
			b	bitstring
			$vk(\hat{A})$	verification key of agent $\hat{A}$
			I	thread-identifier
			$\langle t, \ldots, t \rangle$	tuple
			$\pi_i(t)$	i^{th} component of tuple t
(actions)	α	::=	`send` $\langle \hat{A}, t \rangle$	send a term t to $\hat{A}$
			`receive`	receive a term
			`new`	generate nonce n
			`sign` t	sign term
			`verify` $\langle t, t, vk(\hat{A}) \rangle$	verify a signature
(program)	P	::=	`stop`	terminated program
			$l{:}x \leftarrow \alpha; P$	sequence of actions
(role definition)	R	::=	$r(y) \triangleq [P]_I$	program P with input parameter y

Table 1: Syntax of the Protocol Programming Language

- *Safety Properties.* Finally, since concrete security involves computing the probability of a property being violated as a function of the time available to the adversary, the total runtime of a trace is bounded; traces end when the total time elapsed is greater than the specified time bound. Thus, we cannot talk about liveness properties (i.e., things that will eventually happen) in QPCL.

4 Protocol Programming Language

We now describe the syntax and semantics of the protocol programming language QPCL.

4.1 Protocol Syntax

4.1.1 Syntax: Overview

The protocol programming language is used to represent cryptographic protocols between principals. The language is used to describe the set of roles that comprises a cryptographic protocol. Roles are programs that honest participants of the protocol run instances of. For example, the widely used SSL/TLS protocol has two roles for servers and clients. Each program is a sequence of actions that represent stateful probabilistic computations on bitstrings. Examples of actions are network communication and cryptographic computations such as message signing and verification. A bitstring may denote a principal name, nonce, string constant, signature, or a session identifier. As a matter of convention, the metavariable ($\hat{A}, n$, etc.) that we use to refer to a particular bitstring denotes its intended usage (resp. principal name, nonce, etc.). The syntax of the programming language is summarized in Table 1 and explained below.

Names, sessions and thread identifiers Honest principals can execute multiple instances of different protocol roles. A thread, which is an instance of a protocol role executed by an honest principal, is associated with a *thread identifier* $\langle \hat{A}, \iota \rangle$, a pair containing the principal name $\hat{A}$ and a session identifier ι. If I is a thread identifier, $\pi_1(I)$ is the principal executing the thread.[2] Each thread is assumed to have been assigned an asymmetric signing-verification key pair.[3] All threads of a principal $\hat{A}$ share the same key pair, denoted $(sk(\hat{A}), vk(\hat{A}))$. A signing key $sk(\hat{A})$ is not a term in the programming language, since we do not allow signing keys to appear in messages of a protocol.

Terms, Actions, and Programs The *term language* is composed of variables, bitstrings, verification keys, thread identifiers, tuples and projections.

Actions represent stateful computations on bitstrings that return bitstrings. Examples of actions are nonce generation, signing, verification, and communication (sending and receiving). An *action statement* $x \leftarrow \alpha$ denotes the binding of variable x to the result of the action α. We assume that every bound variable is unique; uniqueness can be ensured by systematically renaming variables. Each action statement has a unique label l. Actions have the form of $a(t)$, where a is a name for the action and t is the input. As a matter of convenience, action labels are not shown explicitly in our examples.

Programs are sequences of action statements. We use the notation $[t/x]P$ to denote substituting term t for variable x in program P (renaming bound variables if necessary to avoid capture).

[2] π_i represents the i^{th} projection.

[3] We do not reason about key exchange in the logic, and assume keys are pre-assigned before the protocol executes.

$$
\begin{array}{lcl}
A \to B & : & m \\
B \to A & : & n, SIG_B\{\!|\text{“}\texttt{Resp}\text{”}, n, m, A|\!\} \\
A \to B & : & SIG_A\{\!|\text{“}\texttt{Init}\text{”}, n, m, B|\!\}
\end{array}
$$

Figure 1: Challenge-response protocol as arrows-and-messages

Role Definitions and Protocols A *role* is a program, instantiations of which are run by honest participants of a protocol. It is defined using a *role definition* of the form of $r(y) \triangleq [P]_I$, which identifies a program P with a name r and the input parameters y for the role. Every role has an additional parameter, denoted by the subscript I, for the thread identifier associated with a particular instance of this role. Every variable in a role is bound, and is either a parameter or assigned to in an action statement. A *protocol* specifies a finite list of roles.

Example Protocol The two roles of the challenge-response protocol depicted informally in Figure 1 are written formally in Figure 2, with $\mathbf{Init_{CR}}$ the initiator role and $\mathbf{Resp_{CR}}$ the responder role. The initiator role $\mathbf{Init_{CR}}$ begins with a *static parameter* $\hat{Y}$, which is the name of the agent that the initiator intends to communicate with. If an agent Alice, for example, uses $\mathbf{Init_{CR}}$ to authenticate an agent Bob using the CR protocol, then Alice executes the instance $\mathbf{Init_{CR}}(\mathbf{Bob})$ of the initiator role obtained by substituting Bob for the static parameter $\hat{Y}$. The value of the static parameters (the choice of the intended responder in this case) is chosen by the adversary before a role is executed.

In words, the actions of a principal executing the role $\mathbf{Init_{CR}}$ are: generate a fresh random number; send a message with the random number to the peer $\hat{Y}$; receive a message with source address $\hat{Y}$; verify that the message contains $\hat{Y}$'s signature over the data in the expected format; and finally, send another message to $\hat{Y}$ with the initiator's signature over the nonce sent in the first message, the nonce received from $\hat{Y}$, and $\hat{Y}$'s identity. The subscript X on the square brackets $[\ \ldots\ ]_X$ enclosing the actions of $\mathbf{Init_{CR}}$ indicate that X is the thread-identifier of the thread executing these actions. X is a pair $\langle \hat{X}, \iota \rangle$ for some principal $\hat{X}$, and thread name ι. A role $[\ \ldots\ ]_X$ may include an action that signs a message with the private key of $\hat{X}$, but not with the private key of $\hat{Y}$ unless $\hat{Y} = \hat{X}$.

To simplify the presentation, we use pattern-matching notation such $(\langle x_1, x_2 \rangle \leftarrow \alpha; P)$ for binding more than one variable in an action statement involving tuples. This should be read as the program with the variables in the pattern replaced with the correct component of the tuple in the rest of the program. For example, the above program should be read as shorthand for $(x \leftarrow \alpha; [\pi_1(x)/x_1, \pi_2(x)/x_2]P)$. We also use _ in place of a variable when the return value from an action is never used.

$\mathbf{Init_{CR}}(\hat{Y}) \triangleq [$

$m \leftarrow$ `new` ;

$_ \leftarrow$ `send` $\langle \hat{Y}, m \rangle$;

$\langle \langle y, s \rangle, N' \rangle \leftarrow$ `receive`;

$_ \leftarrow$ `verify` $\langle s, \langle$"`Resp`"$, y, m, \hat{X} \rangle, vk(\hat{Y}) \rangle$;

$r \leftarrow$ `sign`$\langle$"`Init`"$, y, m, \hat{Y} \rangle$;

$_ \leftarrow$ `send` $\langle \hat{Y}, r \rangle$;

`stop`

$]_X$

$\mathbf{Resp_{CR}}() \triangleq [$

$\langle x, \hat{X} \rangle \leftarrow$ `receive`;

$n \leftarrow$ `new` ;

$r \leftarrow$ `sign`$\langle$"`Resp`"$, n, x, \hat{X} \rangle$;

$_ \leftarrow$ `send` $\langle \hat{X}, \langle n, r \rangle \rangle$;

$\langle t, N'' \rangle \leftarrow$ `receive`;

$_ \leftarrow$ `verify`$\langle t, \langle$"`Init`"$, n, x, \hat{Y} \rangle, vk(\hat{X}) \rangle$;

`stop`

$]_Y$

Figure 2: Roles of the Challenge-response protocol

4.2 Semantics

The semantics of the protocol programming language describe how protocol programs interact in the presence of an adversary who is allowed to execute arbitrary polynomial-time computations.

4.2.1 Semantics: Overview

The semantics describe the interleaved execution of a collection of protocol programs to produce traces. At a given point in time, at most one protocol program in a collection is active, while other programs wait to be executed. Recall that a program is a list of actions. We describe how each action is executed and how actions in different programs are interleaved.

Individual actions are executed by invoking stateful probabilistic computations on bitstrings called *action implementations*. For example, informally, to execute a signature-signing action, the input to the `sign` action is passed to an implementation of digital signature schemes that maintains its own state. Similarly, to execute the sending of a message, the input to the `send` action is sent to the adversary, who is responsible for relaying messages. After the com-

putation implementing the action has finished executing and returns a value, the remainder of the active program continues execution. Certain actions are allowed to make scheduling decisions instead of simply returning values. For example, the receive action is allowed to signal the active program to wait, simulating the adversary's power to make scheduling choices when a program is waiting to receive a message. A special computation, the *scheduler*, is invoked when all threads are waiting.

The adversary controls the scheduler and the send and receive actions, and is allowed to run arbitrary computations when these actions are invoked. In effect, the adversary controls network communication and the scheduling of protocol threads.[4] However, the time for executing every action in a trace is recorded and the total runtime of every trace is bounded by a polynomial of the security parameter, which corresponds to a polynomially bounded adversary.

Action transitions represent the behavior of action implementations, the stateful probabilistic computations that implement actions. Action transitions take as input a state for the action implementation and an input bitstring and return a distribution over updated states and output bitstrings.

Terms such as projections of pairs are resolved down to values representing bitstrings before an action is invoked on such a term. Such stateless computations are represented by the *term-evaluation* relation. Finally, action transitions along with the term-evaluation relation are used in defining *configuration transitions* that describe how individual actions are executed and different protocol programs are interleaved.

4.2.2 Semantics: Details and Examples

We now describe term-evaluation and action transitions, and then define configuration transitions formally.

Term Evaluation and Values Terms represent deterministic, effect-free computations on bitstrings. *Values* are closed terms that cannot be simplified further, namely, bitstrings and tuples of bitstrings. The evaluation of term t to value v with cost c is denoted by the term-evaluation relation $t \Downarrow_c v$. For example, we have that $\pi_1(\langle A, B\rangle) \Downarrow_c A$, for some cost c.

Implementations and Implementation Transitions *Implementations* are stateful probabilistic computations on bitstrings corresponding to each action in the programming language. A special implementation, the scheduler Sr, does not correspond to any action, but is responsible for scheduling when no programs are active.

Implementation transitions are probabilistic transition relations that describe the behavior of action implementations. The observable behavior of

[4]The adversary is also allowed to return garbage values after executing actions that are not parsable as terms, but such values are ignored. Therefore, we make the assumption that all return values of actions are parsable as terms.

implementations includes the cost of computing a particular result and the probability with which the particular result is computed. The following transition relation represents one outcome of executing implementation A on input t:

$$\sigma, A(t) \xmapsto{c,p} \sigma', ra.$$

Here, σ denotes the state maintained by A, c denotes the cost of the outcome, p denotes its probability, σ' denotes the output state and ra denotes the result of the computation (called the *reaction*), which has the following form:

$$ra ::= \texttt{ret}(t) \mid \texttt{abort} \mid \texttt{wait} \mid \texttt{switch}[I].$$

The reaction can either return a value of the expected type ($\texttt{ret}(t)$), signal the active program to abort (`abort`) or wait (`wait`), or force a context switch to the thread with identifier I ($\texttt{switch}[I]$). For example, an implementation of the signature verification action may return if the verification succeeds; otherwise, it can signal the thread to stop. The state maintained by an implementation is treated abstractly by the operational semantics. We have deliberately left the low-level details of states unspecified. The soundness proofs do not depend on the specifics of the implementations. Moreover, adversaries modeled using such an abstract representation are allowed to implement arbitrary data structures to mount efficient attacks on the protocol. To build intuition, we now provide examples of simple instances of implementation transitions.

Consider a deterministic implementation *Sign* of the signature action that keeps a count only of the number of messages it has signed as its state. In this case, the action transition has the following form:

$$n, Sign(m) \xmapsto{c,1} n+1, \texttt{ret}(s).$$

Next, consider an example where an implementation *Verify* of the signature verification action `verify` fails to verify a signature:

$$\langle\rangle, Verify(\langle m, m', vk(N)\rangle) \xmapsto{c,1} \langle\rangle, \texttt{abort}.$$

The implementation *Verify* does not maintain any state; its state is thus denoted $\langle\rangle$.

Now consider an implementation *Send* of the `send` action whose state is a set s of sent messages that corresponds to an adversary who randomly drops half the messages:

$$\begin{array}{lll} s, Send(\langle \hat{A}, m\rangle) & \xmapsto{c,0.5} & s, \texttt{ret}(\langle\rangle) \\ s, Send(\langle \hat{A}, m\rangle) & \xmapsto{c,0.5} & s \cup \{\langle \hat{A}, m\rangle\}, \texttt{ret}(\langle\rangle). \end{array}$$

In the first outcome above, the implementation *Send* keeps the set of sent messages unchanged, effectively dropping the message m. In the second outcome, *Send* updates the state s with the new message. In both outcomes, the null tuple $\langle\rangle$ is returned.

Action	Type	Reactions
`new`	$\texttt{unit} \to \texttt{nonce}$	$\texttt{ret}(t)$
`send`	$\texttt{principal_name} \times \texttt{message} \to \texttt{unit}$	$\texttt{ret}(t)$
`receive`	$\texttt{unit} \to \texttt{principal_name} \times \texttt{message}$	$\texttt{ret}(t), \texttt{wait}$
`sign`	$\texttt{message} \times \texttt{sgn_key} \to \texttt{message}$	$\texttt{ret}(t)$
`verify`	$\texttt{message} \times \texttt{message} \times \texttt{sgn_key} \to \texttt{unit}$	$\texttt{ret}(t), \texttt{abort}$

Table 2: Types and permitted reactions for actions

Consider an implementation *Receive* of the `receive` action that signals the active program to wait when it has no messages to relay:

$$\{\}, \mathit{Receive}(\langle\rangle) \stackrel{c,1}{\longmapsto} \{\}, \texttt{wait}.$$

Finally, $\texttt{switch}[I]$ is used by the scheduler to indicate a switch to the thread with identifier I. For example, a scheduler Sr that manages a list of n threads and picks one at random has the following transition for each k:

$$\{I_1 \cdots I_n\}, Sr \stackrel{c,\frac{1}{n}}{\longmapsto} \{I_1 \cdots I_n\}, \texttt{switch}[I_k].$$

We require that the number of possible transitions from a given state be finite and that the transition probabilities sum to 1. For example, if for the state σ, implementation A, and input t, the possible transitions are

$$\sigma, A(t) \stackrel{c_1,p_1}{\longmapsto} \sigma'_1, ra_1 \quad \ldots \quad \sigma, A(t) \stackrel{c_n,p_n}{\longmapsto} \sigma'_n, ra_n,$$

then we require that $\sum_{i=1}^{n} p_i = 1$ and that for all i, $p_i > 0$.

Setup A setup $\mathcal{S}$ is a mapping $\langle\langle . \rangle\rangle_I$ that associates an action a with an action implementation $\langle\langle a \rangle\rangle_I$ for the thread with identifier I, along with an implementation Sr for the scheduler. We will see below that only the part of the setup that pertains to cryptography is specified by a protocol.

Restrictions on Actions We make well-typing assumptions about the reactions of action implementations. These assumptions are easily discharged, since straightforward checks can identify when responses are not the expected type, and ignore them. Associated with each action is a type and a set of allowed reactions; these are summarized in Table 2 and described here. For instance, according to the table, the only reaction allowed for action $\texttt{send}\langle \hat{A}, t \rangle$ is $\texttt{ret}(\langle\rangle)$; a context switch is not allowed on a send. The implementation for `receive` can respond with $\texttt{ret}(t)$ or `wait`. This means that a context switch can be enforced during a receive action. The fact that only `receive` can respond with a `switch` means that each thread executes its program uninterrupted until it reaches a

`receive` action; at this point, the adversary has the opportunity to schedule other threads.

The action `new` would typically be implemented by a pseudorandom-number generator, and can respond only with $\texttt{ret}(t)$.

The action `sign` returns $\texttt{ret}(t)$; `verify` may respond with a $\texttt{ret}(\langle\rangle)$ on success or `abort` on failure. These actions would typically be implemented by a digital signature scheme.

An implementation may either be specified by the protocol setup or be provided by the adversary. For cryptographic protocols, it may be reasonable to expect that implementations for `send` and `receive` are provided by the adversary, as it controls the communication channel and is responsible for scheduling processes. The implementations for `sign`, `verify`, and `new` are specified by the protocol according to its choice of signature algorithm and pseudorandom-number generator.

Configuration Transitions Finally, configuration transitions describe how protocol programs are executed and interleaved to produce traces. A configuration is a snapshot of the execution of all computation in our system. This snapshot includes each of the protocol programs and the state of each implementation. Formally, a *configuration* $\mathcal{C}$ has the form $E\{\Psi \rhd T\}$. Ψ is a set of waiting threads, where each thread has the form $I : P$, consisting of I, the thread-identifier, and P, the program. T, the thread after the wedge, is the active program being processed. When no program is active, a configuration is denoted $E\{\Psi \rhd \cdot\}$. E, known as the *global state*, is the collection of the states for the action implementations.

Global State The structure of the global state E captures which action implementations are allowed to share information. For example, the adversary is allowed to combine any information that he infers from scheduling or that he receives from the network. However, we require that the pseudorandom-number generators from different threads operate on different states. The global state is a set that has the form $\{h_1 : \sigma_1, \cdots, h_k : \sigma_k\}$, where each σ_i is a local state for some action implementation. The label h_i determines which implementation state σ_i belongs to. Implementations that share state have the same label in the global state. We need to define for each action how the label for that action in the global state is determined. The label of an action a in a thread with identifier I in the global state is denoted $\widetilde{a_I}$. If the implementations of action a in thread I and action b in thread I' operate on the same state, then we want that $\widetilde{a_I} = \widetilde{b_{I'}}$. Otherwise, if the actions operate on separate states, then we want $\widetilde{a_I} \neq \widetilde{b_{I'}}$. The `send` and `receive` actions, which are implemented by the adversary, share state for every thread. So, for all I, we choose $\widetilde{\texttt{send}_I} = \widetilde{\texttt{receive}_I} = \texttt{a}$, for a constant symbol `a`. On the other hand, the `sign` and `verify` actions have separate state for each principal, so we parameterize the label of these actions with a principal, denoting them $\widetilde{\texttt{send}_I} = \widetilde{\texttt{verify}_I} = \texttt{s}(\pi_1(I))$ for a symbol `s` (recall that $\pi_1(I)$

$$\frac{t \Downarrow_{c'} v \qquad \sigma, \langle\langle a \rangle\rangle_I(v) \overset{c,p}{\longmapsto} \sigma', \mathtt{ret}(v')}{E, \widetilde{a_I} : \sigma\{\Psi \rhd I : x \leftarrow a(t); P\} \xrightarrow[(I,a(v),v')]{c+c',p} E, \widetilde{a_I} : \sigma'\{\Psi \rhd I : [v'/x]P\}} (\mathsf{act})$$

$$\frac{t \Downarrow_{c'} v \qquad \sigma, \langle\langle a \rangle\rangle_I(v) \overset{c,p}{\longmapsto} \sigma', \mathtt{abort}}{E, \widetilde{a_I} : \sigma\{\Psi \rhd I : x \leftarrow a(t); P\} \xrightarrow{c+c',p} E, \widetilde{a_I} : \sigma'\{\Psi \rhd \cdot\}} (\mathsf{abort})$$

$$\frac{t \Downarrow_{c'} v \qquad \sigma, \langle\langle a \rangle\rangle_I(v) \overset{c,p}{\longmapsto} \sigma', \mathtt{wait}}{E, \widetilde{a_I} : \sigma\{\Psi \rhd I : x \leftarrow a(t); P\} \xrightarrow{c+c',p} E, \widetilde{a_I} : \sigma'\{\Psi, I : x \leftarrow a(t); P \rhd \cdot\}} (\mathsf{wait})$$

$$\frac{\sigma, Sr \overset{c,p}{\longmapsto} \sigma', \mathtt{switch}[I]}{E, \widetilde{Sr} : \sigma\{\Psi, I : P \rhd \cdot\} \xrightarrow{c,p} E, \widetilde{Sr} : \sigma'\{\Psi \rhd I : P\}} (\mathsf{switch})$$

Table 3: Operational Semantics for the Protocol Language

is the principal corresponding to I). The action `new` has a separate state for each thread, so we denote its label $\widetilde{\mathtt{new}}_I = \mathtt{p}(I)$, for a symbol `p`. Finally, the scheduler Sr is implemented by the adversary, so Sr's label in the global state is `a`, the same as that of `send` and `receive`.

Configuration transitions have the form $\mathcal{C} \xrightarrow{c,p} \mathcal{C}'$, where c and p denote the runtime cost and probability of the transition. Transitions where an action executes to return some value carry a record of the action below the arrow. This record has the form $(I, a(v), v')$, where I is the active thread, $a(v)$ is the action executed and v' is the value returned. (See rule act below.)

We now present rules for the evolution of configurations. These rules are parametric in the transitions for evaluation of actions and terms. Table 3 presents the transition rules for configurations.[5]

The first three rules, act, abort, and wait, describe how a single action in the active program is executed. In each of these rules, $\langle\langle a \rangle\rangle_I$ is the implementation of the action a for thread I, as per the setup. Each rule describes how a statement is executed. First the input to the action is evaluated to a value v; then the implementation $\langle\langle a \rangle\rangle_I$ is run on v to yield a new state σ' and a reaction, and the change in state is reflected in the global state. According to rule act, when $\langle\langle a \rangle\rangle_I$ returns a value v', the active program $x \leftarrow a(t); P$ transitions to $[v'/x]P$, where all occurences of the variable x are replaced with v'. For example, for the signature action `sign`, when the implementation $\langle\langle \mathtt{sign} \rangle\rangle_I$ is called on some message m, whe implementation returns a signature for m. On the other hand, when the implementation responds with `abort`, the current program is stopped.

[5]In the transition rules, the notation $\Psi, I : P$ denotes $\Psi \cup \{I : P\}$; similarly, $E, \widetilde{a_I} : \sigma$ denotes $E \cup \{\widetilde{a_I} : \sigma\}$.

Similarly, when the implementation responds with `wait`, the active program is moved to the set of waiting programs.

The rule switch describes how different programs in a configuration are interleaved. When an implementation responds with `switch`$[I']$, the thread I' is made the active program.

In each rule, the total computation cost and the probability are recorded in the transition. The transitions due to act contain a record of the action executed in the transition and the value returned. Such transitions will be referred to as *labeled transitions*.

Execution Traces A *trace* $\mathcal{L}$ is a sequence of configuration transitions. A trace records the actions in the execution of the protocol, the runtime cost incurred, and its probability. We say that the action label $(I, a(t), t') \in \mathcal{L}$ if $\mathcal{L}$ contains a labeled transition $\mathcal{C} \xrightarrow[(I,a(t),t')]{c,p} \mathcal{C}'$. We also define a function $rt(\mathcal{L})$, which returns the total cost of execution of the trace (i.e., the sum of the runtime costs for each transition), and a function $p(\mathcal{L})$, which returns the probability of the trace (i.e., the product of the probabilities for each transition). The relation $\mathcal{L}_1 \preceq \mathcal{L}_2$ denotes that $\mathcal{L}_1$ is a prefix of $\mathcal{L}_2$.

Execution Trees An execution tree $\mathcal{B}$ rooted at a configuration $\mathcal{C}$ is either the configuration $\mathcal{C}$, or has the form $(\mathcal{C} \xrightarrow{p_1,c_1} \mathcal{B}_1 \mid \cdots \mid \mathcal{C} \xrightarrow{p_k,c_k} \mathcal{B}_k)$, where $\mathcal{C} \xrightarrow{p_1,c_1} \mathcal{C}_1, \cdots \mathcal{C} \xrightarrow{p_k,c_k} \mathcal{C}_k$ are valid transitions from $\mathcal{C}$ and $\mathcal{B}_1, \cdots, \mathcal{B}_k$ are execution trees rooted at $\mathcal{C}_1, \cdots, \mathcal{C}_k$, respectively.

Informally, a *tree* $\mathcal{B}$ is a branching structure of configuration transitions. We write $\mathcal{L} \in \mathcal{B}$ if $\mathcal{L}$ is a path in $\mathcal{B}$. Every path through a tree is a trace. We extend p to finite trees by taking $p(\mathcal{B}) = \sum_{\mathcal{L} \in \mathcal{B}} p(\mathcal{L})$.

Initialization and Termination For a protocol with roles $\{R_1, R_2, \ldots, R_k\}$, the *instance vector* $\vec{n} = \langle n_1, n_2, \ldots, n_k \rangle$ denotes the number of instances of each of the k protocol roles; the *initialization vector* $\vec{\mathbf{t}} = \langle \vec{t_1}, \vec{t_2}, \ldots, \vec{t_k} \rangle$ is a vector of vectors, where $\vec{t_i}$ is a list of values corresponding to the input parameters of all instances of the ith role. Given an initialization vector, the *initial configuration* consists of all the instances of protocol roles running on each thread, where values from the initialization vector are substituted for the corresponding parameters.

Protocol A protocol $\mathcal{Q}$, is a pair $(\vec{R}, \mathcal{S}_\mathcal{Q})$, where $\vec{R}$ is a list of protocol roles, and $\mathcal{S}_\mathcal{Q}$ is a mapping from actions to implementations for actions corresponding to cryptography, namely, `sign`, `verify`, and `new`.

Adversary An adversary $\mathcal{A}$ is a pair $(\vec{\mathbf{t}}, \mathcal{S}_\mathcal{A})$, where $\mathcal{S}_\mathcal{A}$ is the setup for the `send` and `receive` actions and the scheduler, and $\vec{\mathbf{t}}$ is the initialization vector. The remainder of the setup, in particular, the implementations for the signature and pseudorandom-number generation actions, is specified by the protocol.

Given a runtime bound $tb(\eta)$ that is a polynomial of the security parameter η, a trace $\mathcal{L}$ is considered *final* if all threads are terminated or if there exists a transition from the last state such that the total runtime exceeds $tb(\eta)$. A tree $\mathcal{B}$ is considered final if all traces in $\mathcal{B}$ are final.

Definition 1. *(Feasible Traces) For a protocol $\mathcal{Q}$, setup $\mathcal{S}$, instance vector $\vec{n}$, initialization vector $\vec{\mathbf{t}}$, and polynomial $tb(\eta)$ that bounds the running time, $T_{\mathcal{Q}}(\mathcal{S}, \mathbf{t}, \vec{n}, tb, \eta)$ is the final tree obtained starting from the initial configuration. Given an adversary $\mathcal{A}$ that provides implementations $\mathcal{S}_{\mathcal{A}}$ and initial values $\vec{\mathbf{t}}$, and a protocol $\mathcal{Q}$ that provides implementations $\mathcal{S}_{\mathcal{Q}}$, the execution tree is written as $\mathcal{T}_{\mathcal{Q},\mathcal{A}}(\eta, tb, \vec{n}) = T_{\mathcal{Q}}(\mathcal{S}_{\mathcal{A}} \cup \mathcal{S}_{\mathcal{Q}}, \vec{\mathbf{t}}, \vec{n}, tb, \eta)$.*

The function $\mathcal{T}_{\mathcal{Q},\mathcal{A}}$ serves as the model on which QPCL formulas are interpreted.

5 Quantitative Protocol Composition Logic

In this section, we present the syntax, semantics, and proof system for Quantitative Protocol Composition Logic (QPCL). While the logic is similar to Computational PCL [3], the key technical difference is that we can specify and reason about exact security trace properties in QPCL. Specifically, the axioms and proof rules of the logic are annotated with exact bounds, thereby enabling an exact bound to be derived from a formal axiomatic proof of a security property.

QPCL consists of two kinds of formulas: basic formulas that do not involve probability and are interpreted on execution traces; and conditional formulas of the form of $\varphi_1 \rightarrow^{\epsilon} \varphi_2$ that are annotated with a probability function ϵ and are interpreted on execution trees. The probability function ϵ is parameterized by the security parameter, adversary runtime, and number of sessions of the protocol, and specifies, for a particular setting of these parameters, the probability of the conditional implication not holding.

5.1 Syntax

The syntax of formulas is given in Table 4. We summarize the meaning of formulas informally below, with precise semantics presented subsequently.

Basic formulas do not involve probabilities, and are interpreted over execution traces. We use the same syntax for terms and variables in formulas as for terms and variables in the protocol programming language. For every protocol action, there is a corresponding action predicate that asserts that the action has occurred on the trace. For example, $\mathsf{Send}(I, t)$ holds on a trace where the thread with thread-identifier I has sent the term t, while $\mathsf{New}(I, t)$ holds on a trace where the thread I has created the term t using the random nonce generation action. Similarly, $\mathsf{Sign}(I, t, t')$ holds when a signature t for message t' has been generated, and $\mathsf{Verify}(I, t, t', \hat{B})$ holds when t has been verified to be $\hat{B}$'s signature over t'. $\mathsf{Contains}(t, t')$ holds on a trace if t can be derived from t'

Action Predicates:

$$\mathsf{a} \quad ::= \quad \mathsf{Send}(I,t) \mid \mathsf{Receive}(I,t) \mid \mathsf{Sign}(I,s,t) \mid \mathsf{Verify}(I,s,t,\hat{B}) \mid \mathsf{New}(I,t)$$

Basic Formulas:

$$\varphi \quad ::= \quad \mathsf{a} \mid \mathsf{a} < \mathsf{a}' \mid t = t' \mid \mathsf{Start}(T) \mid \mathsf{Contains}(t,\, t') \mid \mathsf{Honest}(\hat{B}) \mid \varphi \wedge \varphi' \mid \neg\varphi \mid \varphi\ [program]_I\ \varphi' \mid \forall x.\varphi$$

Conditional formulas:

$$\psi \quad ::= \quad \varphi \mid (\varphi \rightarrow^{\epsilon} \varphi') \mid \psi \wedge \psi' \mid \neg\psi \mid \forall x.\psi$$

Table 4: Syntax of the logic

using symbolic actions, such as projection and message-recovery from a signature. $\mathsf{Honest}(\hat{B})$ holds if $\hat{B}$ is a principal associated with a protocol thread. We use $\mathsf{Fresh}(I,t)$ as syntactic sugar for the formula $\mathsf{New}(I,t) \wedge \forall t'(\mathsf{Send}(I,t') \Rightarrow \neg\mathsf{Contains}(t,t'))$, which means that the value of t generated by I is "fresh" in the sense that I generated t using a `new` action and did not send out any message containing t. We use $\mathsf{FirstSend}(I,n,m)$ as syntactic sugar for the formula $\mathsf{Contains}(n,m) \wedge \forall t.((\mathsf{New}(I,n) \wedge \mathsf{Send}(I,t) < \mathsf{Send}(I,m)) \Rightarrow \neg\mathsf{Contains}(n,t))$.

The logic includes modal formulas of the form $\varphi_1[program]_I\varphi_2$, where *program* is a fragment of some protocol's program. This formula holds on a trace $\mathcal{L}$ if, roughly speaking, whenever $\mathcal{L}$ can be split into three parts, $\mathcal{L}_1$, $\mathcal{L}_2$, and $\mathcal{L}_3$ (so that $\mathcal{L} = \mathcal{L}_1;\mathcal{L}_2;\mathcal{L}_3$) such that $\mathcal{L}_2$ "matches" *program* and φ_1 holds on trace $\mathcal{L}_1$, then φ_2 holds on $\mathcal{L}_1;\mathcal{L}_2$.

The logic also includes standard connectives and quantifiers of first-order logic. In addition, we have conditional formulas of the form $\varphi \rightarrow^{\epsilon} \varphi'$. As we said in the introduction, ϵ here is a function. Its arguments are η (the security parameter), tb (the adversary's time bound), and $\vec{n}$ (the number of instances of each role of the protocol being analyzed). The formula $\varphi \rightarrow^{\epsilon} \varphi'$ holds if the conditional probability of φ' given φ is greater than or equal to $1 - \epsilon(\eta, tb, \vec{n})$, for every setting of η, tb and $\vec{n}$. Conditional formulas are interpreted over execution trees parameterized by η, tb and $\vec{n}$. While the syntax of QPCL allows arbitrary functions of η, tb and $\vec{n}$ to be used in conditional formulas, we will see in Section 5.3 that formulas provable using QPCL's proofsystem use a very restricted set of functions, which is the set of linear combinations of the bounds on QPCL's axioms.

We abbreviate a special case of the conditional formula of the form $true \rightarrow^{\epsilon} \varphi$ as $B^{\epsilon}\varphi$. Such formulas will prove useful for specifying exact security properties of cryptographic primitives (such as signature schemes) and protocols (such as matching conversation-style authentication properties [12]).

5.2 Semantics

Basic formulas are interpreted over traces, whereas conditional formulas are interpreted over execution trees.

For basic formulas, we define the semantic relation $\mathcal{L}, V \models \varphi$ denoting that basic formula φ holds on trace $\mathcal{L}$, where V is a *valuation* that maps variables to values. The notation $V[x \mapsto v]$ denotes the valuation that is identical to V except that it maps variable x to value v. Recall that all bound variables are assumed to be unique, as uniqueness can be ensured via renaming. The lifting of V from variables to terms is denoted $\hat{V}$. $\hat{V}(t)$ yields a value after replacing each variable in t by $V(x)$. We write $\mathcal{L} \models \varphi$ if, for all V, $\mathcal{L}, V \models \varphi$.

For conditional formulas, we define the semantic relation $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi$. Here, $\mathcal{T}$ is a function that, like ϵ, takes parameters η, tb, and $\vec{n}$; $\mathcal{T}(\eta, tb, \vec{n})$ is an execution tree. We write $\mathcal{T}, V \models \psi$ if $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi$ for all inputs $(\eta, tb, \vec{n})$. Recall from Definition 1 that a protocol $\mathcal{Q}$ and an adversary $\mathcal{A}$ define a model $\mathcal{T}_{\mathcal{Q},\mathcal{A}}$. $\mathcal{T}_{\mathcal{Q},\mathcal{A}}(\eta, tb, \vec{n})$ is the set of traces of protocol $\mathcal{Q}$ generated using adversary $\mathcal{A}$, security parameter η, time bound $tb(\eta)$ on the adversary's running time, and number of instances of each protocol role given by $\vec{n} = \langle n_1, \ldots, n_k \rangle$. For a set Δ of formulas, we define

$\Delta \models_{\mathcal{Q}} \psi$ iff $\mathcal{T}_{\mathcal{Q},\mathcal{A}}, V \models \Delta$ implies $\mathcal{T}_{\mathcal{Q},\mathcal{A}}, V \models \psi$ for all adversaries $\mathcal{A}$, and valuations V.

Definition of $\mathcal{L}, V \models \psi$ We start by giving the semantics of basic formulas. in the definition below, where the different elements of the trace are given by Definition 1.

- $\mathcal{L}, V \models \mathsf{Send}(I, u)$ iff $(\hat{V}(I), \texttt{send}\ \hat{V}(u), \langle\rangle) \in \mathcal{L}$.
- $\mathcal{L}, V \models \mathsf{Receive}(I, u)$ iff $(\hat{V}(I),\ \texttt{receive}\ , \hat{V}(u)) \in \mathcal{L}$.
- $\mathcal{L}, V \models \mathsf{Sign}(I, u, u')$ iff $(\hat{V}(I), \texttt{sign}\ \hat{V}(u'), \hat{V}(u)) \in \mathcal{L}$.
- $\mathcal{L}, V \models \mathsf{Verify}(I, w, u, \hat{B})$ iff $(\hat{V}(I), \texttt{verify}\ \langle \hat{V}(w), \hat{V}(u), vk(\hat{V}(\hat{B}))\rangle, \langle\rangle) \in \mathcal{L}$.
- $\mathcal{L}, V \models \mathsf{New}(I, u)$ iff $(\hat{V}(I), \texttt{new}, \hat{V}(u)) \in \mathcal{L}$.
- $\mathcal{L}, V \models a_1 < a_2$ iff $\mathcal{L}, V \models a_1$, $\mathcal{L}, V \models a_2$, and $\mathtt{a}_1$ appears before $\mathtt{a}_2$ in $\mathcal{L}$, where $\mathtt{a}_1$ and $\mathtt{a}_2$ are the actions in $\mathcal{L}$ corresponding to a_1 and a_2, respectively. (Informally, this predicate models the temporal ordering on actions.)
- $\mathcal{L}, V \models u = v$ if $V(u) = V(v)$, where $=$ is equality on values.
- $\mathcal{L}, V \models \mathsf{Start}(I)$ if no action label in $\mathcal{L}$ has the form $(\hat{V}(I), ., .)$. (Intuitively, this predicate holds on traces in which the thread with identifier I has executed no actions.)

- $\mathcal{L}, V \models \mathsf{Contains}(u, v)$ iff there exists a series of identity function applications, projections, and message-recovery operations from signatures (deriving message from its signature) constructing $\hat{V}(u)$ from $\hat{V}(v)$.
- $\mathcal{L}, V \models \mathsf{Honest}(\hat{A})$ iff $\hat{A}$ is a principal executing a protocol role on $\mathcal{L}$.
- $\mathcal{L}, V \models \theta \wedge \varphi$ iff $\mathcal{L}, V \models \theta$ and $\mathcal{L}, V \models \varphi$.
- $\mathcal{L}, V \models \neg\varphi$ iff $\mathcal{L}, V \not\models \varphi$.

We next define the semantics of modal formulas of the form $\theta[P]_X\varphi$. As we said above, intuitively, $\theta[P]_I\varphi$ is true on a trace $\mathcal{L}$ if, when $\mathcal{L}$ is split into three pieces $\mathcal{L}_1; \mathcal{L}_2; \mathcal{L}_3$, if $\mathcal{L}_2$ "matches" the program $[P]_I$, and if the precondition θ holds for the trace $\mathcal{L}_1$, then the postcondition φ holds for the trace $\mathcal{L}_1; \mathcal{L}_2$. In order to make this intuitive idea precise, we define a notion of *matching* a program $[P]_I$ to a suffix of a trace. Variables in the postcondition φ can be bound to variables declared within P. Matching $[P]_I$ to a trace may place constraints on how these variables are interpreted by a valuation. For example, consider the formula $true[x \leftarrow \mathtt{new}]_I\mathsf{New}(I, x)$, where x is bound within the program. On matching the program to a trace, a concrete value of x for the particular trace is obtained, which is used to evaluate the postcondition $\mathsf{New}(I, x)$.

We define what it means for a trace $\mathcal{L}$ and valuation V to match program $[P]_I$ between prefixes $\mathcal{L}_b$ and $\mathcal{L}_e$ and produce a valuation V', denoted $\mathcal{L}, V \gg [P]_I|\mathcal{L}_b, \mathcal{L}_e, V'$, by induction on $[P]_I$:

- $\mathcal{L}, V \gg [l : x \leftarrow \alpha]_I|\mathcal{L}_1, \mathcal{L}_2, V[x \mapsto t]$ if $\mathcal{L}_2 \preceq \mathcal{L}$ and $\mathcal{L}_2$ ends in a transition $\mathcal{C} \xrightarrow[(\hat{V}(I), \hat{V}(\alpha), t)]{c} \mathcal{C}'$ for some configurations $\mathcal{C}, \mathcal{C}'$.
- $\mathcal{L}, V \gg [l : x \leftarrow \alpha; P]_I|\mathcal{L}_1, \mathcal{L}_2, V'$, if there exist $\mathcal{L}_1'$, $\mathcal{L}_2'$, and V_1 such that (a) $\mathcal{L}, V \gg [l : x \leftarrow \alpha]|\mathcal{L}_1, \mathcal{L}_1', V_1$ and (b) $\mathcal{L}, V_1 \gg [P]_I|\mathcal{L}_2', \mathcal{L}_2, V'$.

The semantics of the modal formula can now be defined as follows:

- $\mathcal{L}, V \models \theta[P]_I\varphi$ iff, for all $\mathcal{L}_1$, $\mathcal{L}_2$, and valuations V', if $\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V'$ and $\mathcal{L}_1, V \models \theta$, then $\mathcal{L}_2, V' \models \varphi$.

Finally, we give semantics to the universal quantifier in the standard way:

- $\mathcal{L}, V \models \forall x.\, \varphi$ iff $\mathcal{L}, V[x \mapsto v] \models \varphi$ holds for all values v. Note that as with the protocol programming language, the set of values is the set of all bitstrings and tuples of bitstrings.

Definition of $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi$ We now give the semantics of $\theta \rightarrow^{\epsilon} \varphi$, where θ and φ are basic formulas. In order to do so, we associate with each basic formula φ a function $[\![\varphi]\!]$ from execution trees to execution trees. Intuitively, for each execution tree $\mathcal{B}$, $[\![\varphi]\!]^V(\mathcal{B})$ is the tree composed of those traces that satisfy φ under valuation V. Since all basic formulas φ are trace properties, $[\![\varphi]\!]^V(\mathcal{B}) = \{\mathcal{L} \in \mathcal{B} |\ \mathcal{L}, V \models \varphi\}$. The formula $\theta \rightarrow^{\epsilon} \varphi$ is intended to express the fact that the probability of φ given θ is within ϵ of 1. But since ϵ is a function of $\eta, tb, \vec{n}$, the semantic relation is parameterized by $\eta, tb, \vec{n}$. We have

- $\mathcal{T}, (\eta, tb, \vec{n}), V \models \varphi$ if $[\![\varphi]\!]^V(\mathcal{T}(\eta, tb, \vec{n})) = \mathcal{T}(\eta, tb, \vec{n})$
- $\mathcal{T}, (\eta, tb, \vec{n}), V \models \theta \rightarrow^{\epsilon} \varphi$ if

$$\frac{p\left(\, [\![\varphi]\!]^V(\mathcal{T}(\eta, tb, \vec{n})) \cap [\![\theta]\!]^V(\mathcal{T}(\eta, tb, \vec{n}))\,\right)}{p\left(\, [\![\theta]\!]^V(\mathcal{T}(\eta, tb, \vec{n}))\,\right)} \;\geq\; 1 - \epsilon(\eta, tb, \vec{n}),$$

  where we take the ratio to be 1 if $p\left(\, [\![\theta]\!]^V(T(\eta, tb, \vec{n}))\,\right) = 0$.

Recall that $B^{\epsilon}\varphi$ is the special case of $\theta \rightarrow^{\epsilon} \varphi$ where $\theta = \mathit{true}$. Also, it can be shown that $p\left(\, \mathcal{T}(\eta, tb, \vec{n})\,\right) = 1$. Thus, we have

- $\mathcal{T}, (\eta, tb, \vec{n}), V \models B^{\epsilon}\varphi$ if

$$p\left(\, [\![\varphi]\!]^V(\mathcal{T}(\eta, tb, \vec{n}))\,\right) \;\geq\; 1 - \epsilon(\eta, tb, \vec{n}).$$

Boolean operations and universal quantification over conditional formulas are defined the same way as for basic formulas.

- $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi_1 \wedge \psi_2$ iff $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi_1$ and $\mathcal{T}, (\eta, tb, \vec{n}), V \models \psi_2$.
- $\mathcal{T}, (\eta, tb, \vec{n}), V \models \neg\psi$ iff $\mathcal{T}, (\eta, tb, \vec{n}), V \not\models \psi$.
- $\mathcal{T}, (\eta, tb, \vec{n}), V \models \forall x.\, \psi$ iff $\mathcal{T}, (\eta, tb, \vec{n}), V[x \mapsto v] \models \psi$ holds for all values v.

Example (Authentication Property) Informally, a matching conversations property states that after an agent executes a particular role, it is known that all the messages sent during the role were actually received by the intended recipient, except possibly the last one. This means that the role of the adversary is limited to that of a wire, faithfully transferring messages to its intended recipients.

In the running example, a matching-conversations form of the authentication property from the point of view of an initiator is formally expressed in QPCL as:

$$\begin{aligned} B^{\epsilon}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \exists \iota.\ & \\ & (\mathsf{Send}(A, m) < \mathsf{Receive}(B, m) \wedge \\ & \mathsf{Receive}(B, m) < \mathsf{Send}(B, \langle y, s\rangle) \wedge \\ & \mathsf{Send}(B, \langle y, s\rangle) < \mathsf{Receive}(A, \langle y, s\rangle))). \end{aligned} \tag{1}$$

In formula (1), B is the thread identifier $\langle \hat{B}, \iota \rangle$. Intuitively, the formula says that, with uncertainty ϵ, after A executes the initiator's $\mathbf{Init}_{CR}$ role for the example challenge-response protocol, the messages between A and responder B occur in order in all but a fraction ϵ of traces of the protocol. Moreover, the proof of the protocol gives us a concrete bound on ϵ that is roughly[6]

$$\epsilon_S^{\mathsf{uf\text{-}cma}}(tb, n_{init} + n_{resp}, \eta) + 2 \cdot 2^{-\eta} + 2\epsilon_P^{\mathsf{prg}}(tb, n_{init} + n_{resp}, \eta).$$

[6]The actual bound involves a few more terms, and is presented in Section 6.

Here, S is the signature scheme and P is the pseudorandom-number generator used in the protocol. Recall from Section 2 that $\epsilon_S^{\mathsf{uf\text{-}cma}}$ is the exact bound on the $\mathsf{uf\text{-}cma}$ security of the signature scheme S and $\epsilon_P^{\mathsf{prg}}$ is the exact bound on the prg security of the pseudorandom-number generator P. Also, n_{init} and n_{resp} are the number of sessions of the initiator and the responder roles, respectively. The overall bound on the security of the protocol is therefore closely related to the security of the underlying cryptographic primitives. An outline of the proof of this formula is given in Section 6; a complete proof is provided in Appendix A.2.

5.3 QPCL Proof System

The proof system for QPCL includes axioms that capture properties of cryptographic primitives (e.g., unforgeability of signatures), a rule for proving invariants of protocol programs (useful for reasoning about properties that capture the behavior of honest principals executing one or more instances of protocol roles), as well as axioms and rules for first-order reasoning about belief formulas. All axioms are annotated with exact bounds. Inference rules compute bounds for the formula in the consequent from the bounds of formulas in the antecedent (e.g., the conjunction of two belief formulas has a bound that is the sum of the bounds for the two formulas). Since a proof of a security property can be visualized as a tree with axioms at the leaves and rules at intermediate nodes, tracking bounds in this manner ensures that the axiomatic proof yields an exact bound for the final property that is proved.

The exact bound in some of the axioms depends on the protocol $\mathcal{Q}$ whose correctness we are trying to prove. As a consequence, the entailment relation $\vdash_{\mathcal{Q}}$ is parameterized by the protocol $\mathcal{Q}$. Axioms that are sound for any protocol are marked by the unparameterized entailment relation $\vdash$.

The exact bound ϵ is a function of the number of instances of the protocol roles $\vec{n}$, the security parameter η, and the running time of the adversary $tb(\eta)$. Axioms related to cryptography introduce error bounds, while other proof rules simply accumulate these bounds. As a result, error bounds of formulas provable in QPCL are linear combinations of the error bounds introduced by the axioms related to cryptography. In the present version of QPCL, error bounds have one of the following forms:

$$\epsilon ::= 0 \mid \epsilon_{\mathbf{VER}} \mid \epsilon_{\mathbf{FS2}} \mid \epsilon_1 + \epsilon_2.$$

Here, $\epsilon_{\mathbf{VER}}$ and $\epsilon_{\mathbf{FS2}}$ are functions that represent the concrete bounds of axioms **VER** and **FS2** discussed below, and $\epsilon_1 + \epsilon_2$ is the function that is pointwise addition of ϵ_1 and ϵ_2.

The main technical result of this section is that the proof system is *sound*, that is, given protocol $\mathcal{Q}$, $\vec{n}$, η, and tb, each provable formula $B^{\epsilon}\varphi$ holds in each of the semantic models defined by $\mathcal{Q}$, $\vec{n}$, η and $tb(\eta)$ with probability at least $1 - \epsilon(\vec{n}, \eta, tb)$.

The soundness proofs for axioms capturing properties of cryptographic primitives are done via exact security reductions; they tightly relate associated

bounds to the bounds for the underlying primitive (e.g., digital signature or pseudorandom-number generator). The exact security reduction proofs proceed by showing that with high probability, whenever a violation of an axiom is detected on a trace, with high probability, a violation of the cryptographic primitive can be constructed. To detect a violation of an axiom on a trace, however, additional computation needs to be performed, which weakens the probability bound on the protocol. Essentially, the reduction proof says that an adversary can break an axiom in time t if an adversary can break the corresponding primitive $t + \delta$. We call δ the *reduction overhead*. Specifying a concrete bound for δ requires implementation details of the reduction such as the hardware on which the reduction is implemented. We elide these details in the paper. Appendix 5.4 contains the reductions for axioms involving cryptography. The exact bound for the reduction overhead for an axiom can be computed for any concrete implementation of the reduction. We are satisfied when the overhead is a small polynomial in the security parameters and number of sessions.

We summarize below some representative axioms and proof rules (additional inference rules and detailed soundness proofs are in Appendix 5.4).

Axioms about cryptographic primitives The axiom below captures the hardness of forging signatures. The natural reading is that if thread Y verified the signature of principal $\hat{B}$, then some thread ι of that principal must have produced the signature. There is one instance of this axiom for each fixed (constant) principal $\hat{B}$.

VER. $\Delta, \mathsf{Honest}(\hat{B}) \vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}}(\mathsf{Verify}(Y, s, m, \hat{B})$
$\Rightarrow \exists \iota.\mathsf{Sign}(\langle \hat{B}, \iota \rangle, s, m) < \mathsf{Verify}(Y, s, m, \hat{B}))$.

Axiom **VER** is not sound in general. It is sound only if the bound $\epsilon_{\mathbf{VER}}$ is chosen appropriately for the protocol $\mathcal{Q}$ (which is why we use the notation $\vdash_{\mathcal{Q}}$). The appropriate choice of the bound is closely related to $\epsilon_S^{\mathsf{uf\text{-}cma}}$, the security bound for the signature scheme. If $\mathcal{Q}$ involves the protocol roles $R_1, \ldots, R_k$, then we define $\epsilon_{\mathbf{VER}}$ as follows:

$$\epsilon_{\mathbf{VER}}(\eta, tb, \vec{n}) = \epsilon_S^{\mathsf{uf\text{-}cma}}(\eta, tb + \delta, q),$$

where $q = \Sigma_{i=1}^{k} (n_i \times q_i)$ is a bound on the total number of signing actions on a trace, n_i is the number of instances of R_i in $\mathcal{Q}$, q_i is the number of signing actions in role R_i, and δ is the reduction overhead for the axiom that can be bounded by a polynomial that is $O(q_i^2 \eta)$.[7]

The next axiom captures the hardness of predicting random nonces generated using a pseudorandom-number generator (PRG). It states that if thread with identifier X generated nonce n and sent it out for the first time in message m, then if n is computable from a message m' received by thread Y, then that receive action must have occurred after the send action by thread X.

[7]The concrete reduction for **VER** is given in Algorithm 1 in Section 5.4. The overhead δ is the concrete runtime for the reduction.

$$\textbf{FS2.}\quad \begin{array}{l}\Delta \vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{FS2}}}(\mathsf{FirstSend}(X,n,m) \wedge \mathsf{Receive}(Y,m') \\ \qquad\qquad \wedge \mathsf{Contains}(n,m') \\ \qquad\qquad\quad \Rightarrow \mathsf{Send}(X,m) < \mathsf{Receive}(Y,m')).\end{array}$$

FS2 is sound for a protocol $\mathcal{Q}$ if it is *nonce-preserving*, which is the case if every program in $\mathcal{Q}$ satisfies the following two (syntactic) conditions:

1. Fresh nonces (nonces not sent out on the network) should not be contained in the input to the `verify` action.

2. Once a message containing a fresh nonce has been signed, no further messages are signed by the program until that message is sent on the network.

As with **VER**, **FS2** is sound only if the bound $\epsilon_{\mathbf{FS2}}$ is chosen appropriately. The appropriate choice of bound is closely related to $\epsilon_P^{\mathsf{prg}}$, the security bound of the pseudorandom-number generator P specified by protocol $\mathcal{Q}$. We define $\epsilon_{\mathbf{FS2}}$ as follows:

$$\epsilon_{\mathbf{FS2}}(\eta, tb, \vec{n}) = \epsilon_P^{\mathsf{prg}}(\eta, tb + \delta, q_n) + \eta \times q \times 2^{-\eta},$$

where q_n is the total number of `new` actions on a trace, $q = n/\eta * q_r$ is an upper bound on the number of nonces that the adversary can send to a protocol thread on a `receive`, n is the length of the longest message that can be received, q_r is the number of receive actions in a trace, and δ is the reduction overhead for the axiom and can be bounded by a polynomial that is $O(q_r\eta^2)$.[8]

Proof rule for protocol invariants To give the proof rules for invariants, we need a little notation. $IS(\mathcal{Q})$ (read as the *initial segments* of $\mathcal{Q}$) denotes the set of all prefixes of each of the roles in $\mathcal{Q}$. Since, a protocol is composed of only a finite number of roles, the proof rule has only a finite number of antecedents.

$$\textbf{HON.}\quad \frac{\Delta \vdash_{\mathcal{Q}} \forall I.(\mathsf{Start}(I) \Rightarrow \varphi) \qquad \forall P \in IS(\mathcal{Q}).\quad \Delta \vdash_{\mathcal{Q}} \forall I.(\mathsf{Start}(I)[P]_I\varphi)}{\Delta \vdash_{\mathcal{Q}} \varphi}$$

Informally, this rule states that if a property φ holds at the beginning of a trace and is preserved by every protocol program, then φ holds at the end of a trace.

Proof rules for modal formulas We have the following rule and axiom for reasoning about modal formulas:

$$\textbf{PC}_{\uparrow}.\quad \frac{\Delta \vdash B^{\epsilon}\varphi}{\Delta \vdash B^{\epsilon}(\theta[P]_I\varphi)}$$

$$\textbf{PC}_{\Rightarrow}.\quad \Delta \vdash ((\theta[P]_X(\varphi \Rightarrow \psi)) \wedge (\theta[P]_X\varphi)) \Rightarrow (\theta[P]_X\psi).$$

[8] Similar to the above, the concrete reduction for **FS2** is given in Algorithm 2 in Section 5.4.

The variant of $\mathbf{PC}_\uparrow$ that does not mention belief—from φ infer $\theta[P]_I\varphi$—is also sound and, indeed, can be derived from $\mathbf{PC}_\uparrow$ and the rules for reasoning about beliefs given below. We seem to need the stronger version to prove results about protocols of interest. We remark that the axiom $(B^{\epsilon_1}(\theta[P]_I(\varphi \Rightarrow \psi)) \wedge B^{\epsilon_2}(\theta[P]_I\varphi)) \Rightarrow B^{\epsilon_1+\epsilon_2}(\theta[P]_I\psi)$, which we actually use in our proof, can be derived easily from $\mathbf{PC}_\Rightarrow$ and the axioms for belief given below.

We have the following rules for Hoare-Style reasoning:

G1. $\dfrac{\Delta \vdash \theta[P]_X\varphi_1 \quad \Delta \vdash \theta[P]_X\varphi_2}{\Delta \vdash \theta[P]_X\varphi_1 \wedge \varphi_2}$.

G2. $\dfrac{\Delta \vdash \theta_1[P]_X\varphi \quad \Delta \vdash \theta_2[P]_X\varphi}{\Delta \vdash \theta_1 \vee \theta_2[P]_X\varphi}$.

G3. $\dfrac{\Delta \vdash \theta \Rightarrow \theta' \quad \Delta \vdash \theta'[P]_X\varphi' \quad \Delta \vdash \varphi' \Rightarrow \varphi}{\Delta \vdash \theta[P]_X\varphi}$.

We also have the following rule for sequential composition:

S1. $\dfrac{\Delta \vdash \varphi_1[P]_X\varphi_2 \quad \Delta \vdash \varphi_2[P']_X\varphi_3}{\Delta \vdash \varphi_1[PP']_X\varphi_3}$.

First-order reasoning about beliefs Our proof system includes a complete proof system for first-order belief logic that allows us to prove formulas of the form $B^\epsilon\varphi$ assuming a set of formulas of the form $B^{\epsilon'}\psi$. The axioms and inference rules are just specializations of more general axioms and rules given by Halpern [1]. Here is (a slightly simplified version of) the proof system[9]:

B1. $\dfrac{\Delta \vdash \varphi \Rightarrow \varphi'}{\Delta \vdash B^\epsilon\varphi \Rightarrow B^\epsilon\varphi'}$.

B2. $\Delta \vdash B^0(\mathit{true})$

B3. $\Delta \vdash B^\epsilon\varphi \wedge B^{\epsilon'}\varphi' \Rightarrow B^{\epsilon+\epsilon'}(\varphi \wedge \varphi')$

The next axiom is sound in our setting, where all traces get positive probability. This means if φ holds with probability 1, then it must hold on all traces, since no trace has probability 0.

B4. $\Delta \vdash B^0(\varphi) \Rightarrow \varphi$.

We include all first order tautologies as axioms in our proof system.

5.4 Soundness

We now state our soundness theorem. $\vdash_{\mathcal{Q}} \psi$ denotes that ψ is provable using instances of the axiom and proof rule schemas for protocol $\mathcal{Q}$ (that is, all rules of the form $\vdash_{\mathcal{Q}} \ldots$ or $\vdash \ldots$).

Theorem 1. *(Soundness) If* $\vdash_{\mathcal{Q}} \psi$ *then* $\Delta \models_{\mathcal{Q}} \psi$.

[9]The more general versions of axioms B1-3 are presented in Appendix A.3

We first present definitions and lemmas required to prove the soundness of QPCL's proofsystem, and then prove the soundness of the axioms about cryptographic primitives as illustrative cases. Appendix A.1 contains proofs of other axioms.

Each property that is provable using QPCL's proof system is a safety property in the following sense.

Definition 2. *A basic formula φ is called a* safety property *if $\mathcal{L}, V \not\models \varphi$ implies that for all traces $\mathcal{L}'$ such that $\mathcal{L} \preceq \mathcal{L}'$, $\mathcal{L}', V$ and $\not\models \varphi$.*

Safety ensures that properties which are false on a trace continue to be false. This allows us the bound the probability of a formula being false in the middle of an execution, by the probability at the end of the execution.

Since safety applies only to trace properties and not to formulas with belief operators, we convert probabilistic assertions to basic formulas using an erasure operation $(\cdot)^\circ$.

Definition 3 (Erasure)**.** *The erasure operation $(\psi)^\circ$ is defined by induction on ψ as follows:*

- $(B^\epsilon(\varphi))^\circ = \varphi$.
- $(\psi_1 \wedge \psi_2)^\circ = (\psi_1)^\circ \wedge (\psi_2)^\circ$.
- $(\neg\psi)^\circ = \neg(\psi)^\circ$.
- $(\forall x.\psi)^\circ = \forall x.(\psi)^\circ$.
- $(\varphi)^\circ = \varphi$.

Lemma 1 (Safety)**.** *If $\vdash_{\mathcal{Q}} \psi$, then $(\psi)^\circ$ is a safety property.*

Proof. By induction on the derivation of $\vdash_{\mathcal{Q}} \psi$. We consider three cases here, and leave the remaining cases to the reader:

B1. Assume by the induction hypothesis that $\varphi \Rightarrow \varphi'$ is a safety property. Then it is immediate that $(B^\epsilon(\varphi) \Rightarrow B^\epsilon(\varphi'))^\circ$ is a safety property.

PC$_\uparrow$. Assuming, by the induction hypothesis, that φ is a safety property, we must show that $\theta[P]_I\varphi$ is a safety property. Assume that for some $\mathcal{L}$, V, it is the case that $\mathcal{L}, V \not\models \theta[P]_I\varphi$. By the semantics of the modal formula, we know that there exists $\mathcal{L}_1$, $\mathcal{L}_2$, V' such that $(\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V')$, $\mathcal{L}_1, V \models \theta$, and $\mathcal{L}_2, V' \not\models \varphi$. For all traces $\mathcal{L}'$ such that $\mathcal{L} \preceq \mathcal{L}'$, we must have that $\mathcal{L}', V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V'$. Therefore, $\mathcal{L}', V \not\models \theta[P]_I\varphi$.

FS2. We need to show that $\varphi_{\mathbf{FS2}} = \mathsf{FirstSend}(X, n, m) \wedge \mathsf{Receive}(Y, m') \Rightarrow \mathsf{Send}(X, m) < \mathsf{Receive}(Y, m')$ is a safety property. Assume that for some trace $\mathcal{L}$, V, it is the case that $\mathcal{L}, V \not\models \varphi_{\mathbf{FS2}}$. Therefore, it must be the case that $(\hat{V}(Y), \texttt{receive}, \hat{V}(m')) \in \mathcal{L}$ and that $(\hat{V}(X), \texttt{sent}\ \hat{V}(m), \langle\rangle)$ does not appear in $\mathcal{L}$ before $(\hat{V}(Y), \texttt{receive}, \hat{V}(m')) \in \mathcal{L}$. Therefore, the same must be true for all $\mathcal{L}'$ such that $\mathcal{L} \preceq \mathcal{L}'$.

□

Soundness of the axioms about cryptographic primitives We prove axioms about cryptographic primitives sound by reducing the validity of the axiom to the security of the primitive. Such proofs require augmenting the implementations of primitives with a program that monitors the validity of the axiom on the trace. When a violation of the axiom is detected, it can be transformed into a violation of the primitive's security. As every attack on the axiom can be transformed in to an attack on the primitive, the probability of the primitive being broken is a bound on the probability that the axiom is false. However, we need to account for the additional time required to monitor the validity of the axiom while computing the probability of attack on the primitive.

In order to prove the soundness of **VER**, we need to provide a formal definition of execution tree, so that we can do an induction on the structure of the execution tree.

We define a prefix relation on trees. Intuitively, a tree $\mathcal{B}$ is a prefix of $\mathcal{B}'$ if every trace in $\mathcal{B}$ is a prefix of $\mathcal{B}'$.

Definition 4 (Tree Prefix)**.** *A tree $\mathcal{B}$ is a* prefix *of $\mathcal{B}'$, denoted $\mathcal{B} \preceq \mathcal{B}'$, if one of the following conditions hold:*

1. $\mathcal{B} = \mathcal{B}'$,
2. $\mathcal{B}$ *is a configuration* $\mathcal{C}$ *and* $\mathcal{B}' = (\mathcal{C} \to \mathcal{B}_1 \mid \cdots \mid \mathcal{C} \to \mathcal{B}_k)$[10],
3. $\mathcal{B} = (\mathcal{C} \to \mathcal{B}_1 \mid \cdots \mid \mathcal{C} \to \mathcal{B}_k)$, $\mathcal{B}' = (\mathcal{C} \to \mathcal{B}'_1 \mid \cdots \mid \mathcal{C} \to \mathcal{B}'_k)$, *and* $\mathcal{B}_1 \preceq \mathcal{B}'_1 \cdots \mathcal{B}_k \preceq \mathcal{B}'_k$.

We now present an equivalent definition of the probability of a formula interpreted on a tree. The function $\Pr(\varphi, \mathcal{L}, \mathcal{B}, V)$ denotes the probability of a formula φ in the subtree $\mathcal{B}$ that follows after trace $\mathcal{L}$, with valuations V. Figure 3 shows a trace $\mathcal{L}$ followed by a subtree $B_{\mathcal{L}}$ of $\mathcal{B}$. Since QPCL formulas depend on an entire trace, it is not enough to define the probability of a formula with respect to a subtree, but to include the history of events that ocurred before the subtree as illustrated in Figure 3. In the definitions below, $\mathcal{L} \to \mathcal{C}$ denotes the the trace obtained by extending $\mathcal{L}$ with the transition $\mathcal{C}' \to \mathcal{C}$, where $\mathcal{C}'$ is the last configuration in $\mathcal{L}$. The empty trace is denoted by $\cdot$, and $\cdot \to \mathcal{C}$ is the trace $\mathcal{C}$, containing only a single configuration and no transitions.

Definition 5 (Probability)**.** $\Pr(\varphi, \mathcal{L}, \mathcal{B}, V)$ *is defined inductively on $\mathcal{B}$ as follows:*

1. $\Pr(\varphi, \mathcal{L}, \mathcal{C}, V) = 1$ *if* $\mathcal{L} \to \mathcal{C}, V \models \varphi$,
2. $\Pr(\varphi, \mathcal{L}, \mathcal{C}, V) = 0$ *if* $\mathcal{L} \to \mathcal{C}, V \not\models \varphi$,
3. $\Pr(\varphi, \mathcal{L}, \mathcal{B}, V) = \sum_{i=0}^{k} p_k \Pr(\varphi, \mathcal{L} \to \mathcal{C}, \mathcal{B}_i, V)$ *if* $\mathcal{B} = (\mathcal{C} \xrightarrow{p_1, c_1} \mathcal{B}_1 \mid \cdots \mid \mathcal{C} \xrightarrow{p_k, c_k} \mathcal{B}_k)$.

[10] Cost and probability labels are elided here for brevity.

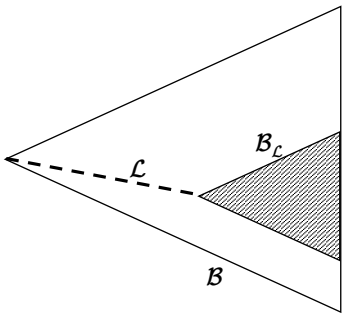

Figure 3: An Execution Tree with a trace $\mathcal{L}$ and subtree $\mathcal{B}_{\mathcal{L}}$

The probability of a formula φ interpreted on a tree $\mathcal{B}$ preceded by the empty trace, and using valuation V is $\Pr(\varphi, \cdot, \mathcal{B}, V)$. The following lemma states that the two definitions of probability are equivalent.

Lemma 2. *For all valuations V, $p\left([[\varphi]]^V(\mathcal{B})\right) = \Pr(\varphi, \cdot, \mathcal{B}, V)$.*

We also state a lemma that relates the probability of a safety property on a tree $\mathcal{B}$ to its probability on a prefix of $\mathcal{B}$.

Lemma 3. *If φ is a safety property and $\mathcal{B} \preceq \mathcal{B}'$, then for all V, $p\left([[\varphi]]^V(\mathcal{B})\right) \geq p\left([[\varphi]]^V(\mathcal{B}')\right)$.*

In our proofs of the soundness of the axioms about cryptographic primitives, we replace the implementations of some actions with different implementations that preserve input-output behavior, in a sense made precise by the following definition.

Definition 6 (Simulation)**.** *An implementation $\bar{A}$* simulates *an implementation A with the relation $\sim$ on $S_A \times S_{\bar{A}}$ (where S_B denotes the set of states of implementation B), if (a) $\sigma_A^{init} \sim \sigma_{\bar{A}}^{init}$, where σ_A^{init} and $\sigma_{\bar{A}}^{init}$ are the initial states of A and $\bar{A}$ respectively and (b) if $\sigma_A \sim \sigma_{\bar{A}}$, then for all transitions $\sigma_A, A(t) \xmapsto[p]{c} \sigma'_A, ra$ of A, there exists a transition $\sigma_{\bar{A}}, \bar{A}(t) \xmapsto[p]{c'} \sigma'_{\bar{A}}, ra$ of $\bar{A}$ such that $\sigma'_A \sim \sigma'_{\bar{A}}$. If there exists a b such that, for all transitions of A, if c is the cost of the transition and c' is the cost of the corresponding transitions of $\bar{A}$, we have $c \leq c' \leq c + b$, then $\bar{A}$ simulates A* with time bound b.

Intuitively, if $\bar{A}$ simulates A with the relation $\sim$, then starting from $\sim$-related states, the implementations A and $\bar{A}$ exhibit the same input-output behavior, but have possibly different costs.

We wish to lift simulations over action implementations to simulations over setups, which are collections of action implementations. However, for a setup $\mathcal{S}'$ to simulate $\mathcal{S}$, it is not sufficient for the corresponding implementations in $\mathcal{S}'$ and $\mathcal{S}$ to be simulations. Recall that different implementations may share state,

and that state sharing is governed by the label $\widetilde{a}_I$ of an action a in the thread with identifier I in the global state. Therefore, we require implementations that share state to respect each other's simulation relations.

Definition 7 (Setup Simulation)**.** *A setup $\mathcal{S}'$ simulates a setup $\mathcal{S}$ if there exists a set $\{\sim_{h_1}, \cdots, \sim_{h_n}\}$ of relations, indexed by global state labels $\{h_1, \cdots, h_n\}$, and each $\langle\langle a \rangle\rangle_I$ in $\mathcal{S}'$ simulates the corresponding implementation in $\mathcal{S}$ with relation $\sim_h$, where $h = \widetilde{a}_I$.*

We now show that if two setups $\mathcal{S}$ and $\mathcal{S}'$ differ only in their action implementations, and the action implementations in $\mathcal{S}'$ simulate the corresponding implementations in $\mathcal{S}$, then for any safety property φ, and assignment V, the probability of φ in the tree generated by $\mathcal{S}$ is bounded by the probability of φ in the tree generated by $\mathcal{S}'$.

Lemma 4. *If setup $\mathcal{S}'$ simulates setup $\mathcal{S}$, each implementation $\langle\langle a \rangle\rangle_I$ in $\mathcal{S}'$ simulates the corresponding implementation in $\mathcal{S}$ with time bound b_{a_I}, and φ is a safety property, then for all V, $p\left([[\varphi]]^V(\mathcal{T}(\mathcal{S}, \mathbf{t}, \vec{n}, tb, \eta))\right) \geq p\left([[\varphi]]^V(\mathcal{T}(\mathcal{S}', \mathbf{t}, \vec{n}, tb + \Sigma_{a_I} n_{a_I} b_{a_I}, \eta))\right)$, where n_{a_I} is the number of times implementation a_I is called.*

Proof. We first show by induction on $\mathcal{T}(\mathcal{S}, \mathbf{t}, \vec{n}, tb, \eta)$ that $\mathcal{T}(\mathcal{S}, \mathbf{t}, \vec{n}, tb, \eta) \preceq \mathcal{T}(\mathcal{S}', \mathbf{t}, \vec{n}, tb + \Sigma_{a_I} n_{a_I} b_{a_I}, \eta)$. Essentially, we can show that each trace using the implementations in $\mathcal{S}'$ makes transitions identical to the trace using the implementations in $\mathcal{S}$, and takes no more time than the extra time allotted ($\Sigma_{a_I} n_{a_I} b_{a_I}$). We then apply Lemma 3. □

In the theorem above, we call the extra time allocated for the simulation to execute ($\Sigma_{a_I} n_{a_I} b_{a_I}$) the *reduction overhead*.

Lemma 5. *The inference rule* **VER** *is sound.*

Proof. Let $\varphi_{\mathbf{VER}}$ abbreviate the formula $\mathsf{Verify}(Y, s, m, \hat{B}) \Rightarrow \exists \iota.\mathsf{Sign}(\langle \hat{B}, \iota \rangle, s, m)$. We must show that $\models_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}} \varphi$, where $\epsilon_{\mathbf{VER}}$ is defined as follows:

$$\epsilon_{\mathbf{VER}}(\eta, tb, \vec{n}) = \epsilon_S^{\mathsf{uf\text{-}cma}}(\eta, tb + \delta, q),$$

where, $\epsilon_S^{\mathsf{uf\text{-}cma}}$ is the concrete security bound of the signature scheme S specified by the protocol $\mathcal{Q}$, $q = \Sigma_{i=1}^{k} (n_i \times q_i)$ is a bound on the total number of signing actions on a trace, n_i is the number of instances of R_i in $\mathcal{Q}$, q_i is the number of signing actions in role R_i, and δ is the reduction overhead for the simulation where the implementations of actions `sign` and `verify` are replaced respectively by the implementations `sign`′ and `verify`′ described in Algorithm 1.

Given $\mathcal{A}, \eta, tb, \vec{n}$, and V, we want to show that

$$p\left([\![\neg\varphi_{\mathbf{VER}}]\!]^V(\mathcal{T}_{\mathcal{Q},\mathcal{A}}(\vec{n}, tb, \eta))\right) \leq \epsilon_{\mathbf{VER}}(\eta, tb, \vec{n}).$$

As a first step, we replace the implementations of `sign` and `verify` in $\mathcal{Q}$ with the implementations shown in Algorithm 1. The implementation `sign`′ keeps track of the messages signed; `verify`′ checks whether a verified message has

ALGORITHM 1: Implementations $\texttt{sign}'$ and $\texttt{verify}'$ augment `sign` and `verify` to monitor **VER**

Implementation $\texttt{sign}'((\sigma, M), (m, k))$

$(\sigma', r) \leftarrow \texttt{sign}(\sigma, (m, k))$;
return $(\sigma', M \cup \{m\}), r$;

Implementation $\texttt{verify}'((\sigma, M), (m, s, k))$

$(\sigma', r) \leftarrow \texttt{verify}(\sigma, (m, s, k))$;
$bad \leftarrow (m \notin M) \wedge (r \neq \texttt{abort})$;
return $(\sigma', M), r$;

been signed before. The variable *bad* in $\texttt{verify}'$ is set to true when a message is verified that has not been signed before. Therefore, *bad* is set to true precisely when $\neg\varphi_{\mathbf{VER}}$ is true on a trace. Note that the implementations $\texttt{sign}'$ and $\texttt{verify}'$ simulate `sign` and `verify`. The simulation relation is $\sigma \sim (\sigma, M)$ for all σ and M, and the initial state for $\texttt{sign}'$ and $\texttt{verify}'$ is $(\sigma^{init}, \emptyset)$, where σ^{init} is the initial state for the implementation of `sign` and `verify`. Therefore, by Lemma 4, we have that

$$p\left([\![\neg\varphi_{\mathbf{VER}}]\!]^V(\mathcal{T}_{\mathcal{Q},\mathcal{A}}(\vec{n}, tb, \eta))\right) \leq p\left([\![\neg\varphi_{\mathbf{VER}}]\!]^V(\mathcal{T}_{\mathcal{Q}',\mathcal{A}}(\vec{n}, tb+\delta, \eta))\right).$$

Let $\mathcal{Q}'$ be the protocol that results by replacing the implementations `sign` and `verify` in $\mathcal{Q}$ with $\texttt{sign}'$ and $\texttt{verify}'$. Let δ be the additional runtime required for the reduction in Algorithm 1; δ can be bounded by a polynomial that is $O((q_i + q_j)^2\eta)$, where q_i and q_j are the number of signature and verification actions in a role, assuming some quadratic implementation for the set operations in Algorithm 1. Recall from Section 2 that the uf-cma security condition is violated when the adversary is able to create a signature of a message that has not been signed before. Therefore, when *bad* is set to true, the uf-cma security of the scheme S has been violated. Consequently, the probability that *bad* is set to true can be bounded by $\epsilon_S^{\mathsf{uf\text{-}cma}}(tb+\delta, q, \eta)$, where q is the number of signing actions and is bounded by the number of signing operations possible in running the protocol, that is, $q = \Sigma_{i=1}^{k}\,(n_i \times q_i)$ and q_i is the number of signing actions in role ρ_i of protocol $Q = r_1, \ldots, r_k$. Since *bad* is set to true when $\neg\varphi_{\mathbf{VER}}$ is true on a trace, we can show that

$$p\left([\![\neg\varphi_{\mathbf{VER}}]\!]^V(\mathcal{T}_{\mathcal{Q}',\mathcal{A}}(\vec{n}, tb+\delta, \eta))\right) \leq \epsilon_S^{\mathsf{uf\text{-}cma}}(tb+\delta, q, \eta) = \epsilon_{\mathbf{VER}}(\eta, tb, \vec{n}).$$

□

Lemma 6. *The axiom* **FS2** *is sound.*

Proof. Let $\varphi_{\mathbf{FS2}}$ abbreviate $\mathsf{FirstSend}(X, n, m) \wedge \mathsf{Receive}(Y, m') \wedge \mathsf{Contains}(n, m') \Rightarrow \mathsf{Send}(X, m) < \mathsf{Receive}(Y, m')$. We must show that $\models_{\mathcal{Q}} B^{\epsilon_{\mathbf{FS2}}}\varphi_{\mathbf{FS2}}$, where $\epsilon_{\mathbf{FS2}}$ is defined as follows:

$$\epsilon_{\mathbf{FS2}}(\eta, tb, \vec{n}) = \epsilon_P^{\mathsf{prg}}(\eta, tb+\delta, q_n) + \eta \times q \times 2^{-\eta},$$

where q_n is the total number of `new` actions on a trace, $q = l/\eta * q_r$ is an upper bound on the number of nonces that the adversary can send to a protocol thread on a `receive`, l is the length of the longest message that can be received, q_r is the number of new, send, and receive actions in a trace, and δ is the reduction overhead for the simulation where the implementations of actions `new`, `send` and `receive` are replaced respectively by the implementations `new′`, `send′` and `receive′` described in Algorithm 2.

ALGORITHM 2: Implementations `new′`, `send′`, and `recv′` augment `new`, `send` and `recv` repectively to monitor **FS2**

Implementation `new'`$((\sigma, M), ())$

1 $(\sigma', \texttt{ret}(n)) \leftarrow \texttt{new}(\sigma, ())$;
2 **return** $(\sigma', M \cup \{n\}), \texttt{ret}(n)$;

Implementation `send'`$((\sigma, M), (m, k))$

1 $(\sigma', r) \leftarrow \texttt{send}(\sigma, (m, k))$;
2 **return** $(\sigma', M - \texttt{nonces}(m)), r$;

Implementation `recv'`$((\sigma, M), ())$

1 $(\sigma', r) \leftarrow \texttt{recv}(\sigma, ())$;
2 **if** $r = \texttt{ret}(m)$ **then** $bad \leftarrow \exists n \in \texttt{nonces}(m) \wedge n \in M$;
3 **return** $(\sigma', M), r$;

Given $\mathcal{A}, \eta, tb$, and $\vec{n}$, we must show that

$$p\left([\![\neg\varphi_{\mathbf{FS2}}]\!]^V(\mathcal{T}_{\mathcal{Q},\mathcal{A}}(\vec{n}, tb, \eta))\right) \leq \epsilon_{\mathbf{FS2}}(\eta, tb, \vec{n}).$$

Let $\mathcal{Q}'$ be the protocol that results from replacing the implementations of `new`, `send`, and `receive` with the implementations shown in Algorithm 2. The set M records the nonces that have been generated but not been sent out. The implementation `new′` records the nonces that have been generated. When a message is sent with a nonce using `send′`, it is removed from M. The variable *bad* is set to true when there is a nonce in M that has been received. Essentially, *bad* is set to true on a trace exactly when $\varphi_{\mathbf{FS2}}$ is false.

Consider the following two experiments:

- $\mathbf{Exp}^{\mathrm{Random}}$: Run protocol $\mathcal{Q}'$, with `new` implemented by a random-number generator with uniformly random distibution. If *bad* is set to true output 1 else 0.
- $\mathbf{Exp}^{P}$: Run protocol $\mathcal{Q}'$, with `new` implemented by the pseudorandom-number generator P specified by the protocol. If *bad* is set to true output 1 else 0.

Recall from Section 2 that the probability of an adversary being able to distinguish a uniform random number generator from the output of a pseudorandom-number generator S after seeing q samples and in time t is bounded by $\epsilon_P^{\mathsf{prg}}(\eta, t, q)$. Therefore, we have the following:

$$\mathbf{Pr}(\mathbf{Exp}^{P} = 1) - \mathbf{Pr}(\mathbf{Exp}^{\mathrm{Random}} = 1) \leq \epsilon_P^{\mathsf{prg}}(\eta, tb + \delta, q),$$

where $q_n = \Sigma_{i=1}^{k} (n_i \times q_i)$, where q_i is the number of `new` actions in role i and n_i is the number of instances of role r_i in protocol roles $\{r_1, \cdots, r_k\}$. The additional runtime δ can be bounded by a polynomial that is $O(q_r^2 \eta)$, where q_r is the total number of new, send, and receive actions in the protocol, assuming some quadratic implementation of the set operations in Algorithm 2. The total runtime for ϵ^{prg} is $tb + \delta$ to account for the additional runtime for the reduction.

As *bad* is set to true exactly when $\neg\varphi_{\mathbf{FS2}}$ is true, as with the proof for **VER**, we can use Lemma 4 to get the following bound for $\varphi_{\mathbf{FS2}}$:

$$p\left([\![\neg\varphi_{\mathbf{FS2}}]\!]^V \mathcal{T}_{\mathcal{Q},\mathcal{A}}(\vec{n}, tb, \eta)\right) \leq \mathbf{Pr}(\mathbf{Exp}^P(A) = 1).$$

We now only need to derive a bound for $\mathbf{Pr}(\mathbf{Exp}^{\mathrm{Random}} = 1)$, since $\mathbf{Pr}(\mathbf{Exp}^P = 1) \leq \mathbf{Pr}(\mathbf{Exp}^{\mathrm{Random}}(A) = 1) + \epsilon^{\mathsf{prg}}(\eta, tb + \delta, q)$. When we use a true random number generator, at each guess, the adversary can guess the random number generated with probability only $2^{-\eta}$. Therefore, if the adversary has q guesses, by the union bound, we can show that the total probability of *bad* being true is $q2^{-\eta}$. The adversary can attempt to guess the nonce at each `receive` action. However, each message received can contain multiple nonces. If the longest message in the protocol is of length l, then each message received can contain at most l/η nonces of length η, and so a bound on q, the total number of guesses available to the adversary, is $q_r l/\eta$. Finally, we have that

$$\mathbf{Pr}(\mathbf{Exp}^{\mathrm{Random}} = 1) \leq \frac{q_r l}{\eta 2^{\eta}}.$$

Therefore, we have that

$$p\left([\![\neg\varphi_{\mathbf{FS2}}]\!]^V \mathcal{T}_{\mathcal{Q},\mathcal{A}}(\vec{n}, tb, \eta)\right) \leq \mathbf{Pr}(\mathbf{Exp}^P = 1) \leq \epsilon^{\mathsf{prg}}(\eta, tb+\delta, q) + \frac{q_r l}{\eta 2^{\eta}} = \epsilon_{\mathbf{FS2}}(\eta, tb, \vec{n}).$$

□

6 Proof Outline for the Example Protocol

A formal proof of the weak authentication property for the initiator guaranteed by executing the CR protocol is presented in Appendix A.2. We explain below the outline of the proof and discuss how the axioms and rules described in the previous section are used. The formal proof proceeds roughly as follows.

First, we assert what actions were executed by the thread with identifier A in the initiator role. Specifically, in this part of the proof, we prove that A has received and verified $\hat{B}$'s signature. We then use the fact that the signatures of honest parties are unforgeable (axiom **VER**) and first-order reasoning about beliefs to conclude that some thread with name ι of $\hat{B}$ ($B = \langle \hat{B}, \iota \rangle$ below) must have produced this signature. This step introduces a bound of $\epsilon_{\mathbf{VER}}$, and we obtain the following formula:

$$\mathsf{Honest}(\hat{B}) \vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \ \mathsf{Sign}(B, s, \langle \text{“}\texttt{Resp}\text{”}, y, m, \hat{A}\rangle)). \quad (2)$$

Second, we use the honesty rule **HON** to infer that whenever any thread (and by instantiation using the rules of first-order belief logic, any the thread with identifier B of agent $\hat{B}$) generates a signature of this form, he must have previously received the first message from $\hat{A}$ and then sent the second message; moreover, the second message was in fact the first send involving the random nonce y:

$$\begin{aligned}\vdash_{\mathcal{Q}} \quad & B^0(\mathsf{Sign}(B, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle) \Rightarrow (\mathsf{Receive}(B, m) \\ & \quad \wedge (\mathsf{Send}(B, \langle y, s\rangle) \Rightarrow \\ & \qquad \mathsf{Receive}(B, m) < \mathsf{Send}(B, \langle y, s\rangle)) \\ & \quad \wedge \mathsf{FirstSend}(B, \langle y, s\rangle))).\end{aligned} \tag{3}$$

Combining (2) and (3) using rules for first-order reasoning about beliefs, we conclude that the following formula, which orders the actions performed by B, holds:

$$\begin{aligned}\vdash_{\mathcal{Q}} \quad & B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \\ & \qquad \mathsf{Receive}(B, m) < \mathsf{Send}(B, \langle y, s\rangle)).\end{aligned} \tag{4}$$

Next we want to prove that B received the first message after A sent it. We use the fact that since the thread with identifier A generated the random nonce m and sent it out for the first time in the first message of the protocol, then if the thread with identifier B received a message from which m is computable, he must have done so after A sent out the first message. This step uses the **FS2** axiom, and therefore the bound associated with the underlying PRG applies:

$$\begin{aligned}\vdash_{\mathcal{Q}} \quad & B^{\epsilon_{\mathbf{FS2}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \\ & \quad ((\mathsf{FirstSend}(A, m, m) \wedge \mathsf{Receive}(B, m) \wedge \mathsf{Contains}(m, m)) \\ & \qquad \Rightarrow \mathsf{Send}(A, m) < \mathsf{Receive}(B, m))).\end{aligned} \tag{5}$$

Using the fact that A's first send had the form in the premise of the implication in (5) and that B received the corresponding message from (3), from B3, (4), and (5), we infer desired ordering of actions by A and B. Note that since two belief formulas are combined, the bounds add up:

$$\begin{aligned}\mathsf{Honest}(\hat{B}) \vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}+\epsilon_{\mathbf{FS2}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \\ \mathsf{Send}(A, m) < \mathsf{Receive}(B, m)).\end{aligned}$$

The rest of the proof uses similar ideas (including another application of **FS2**) to order other actions performed by threads A and B to prove a matching-conversation like property of the protocol with a bound that is tightly related to the security of the underlying digital signature scheme and pseudorandom-number generator:

$$\begin{aligned}\mathsf{Honest}(\hat{B}) \vdash_{\mathcal{Q}} \quad & B^{\epsilon_{\mathbf{VER}}+2\cdot\epsilon_{\mathbf{FS2}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \ \exists \iota. \\ & \quad \mathsf{Send}(A, m) < \mathsf{Receive}(B, m) \wedge \\ & \quad \mathsf{Receive}(B, m) < \mathsf{Send}(B, \langle y, s\rangle) \wedge \\ & \quad \mathsf{Send}(B, \langle y, s\rangle) < \mathsf{Receive}(A, \langle y, s\rangle)),\end{aligned}$$

where $B = \langle \hat{B}, \iota \rangle$. The bound on the security guarantee is, $\epsilon = \epsilon_{\mathbf{VER}} + 2\epsilon_{\mathbf{FS2}}$. Therefore $\epsilon(\eta, tb, \vec{n}) = \epsilon_S^{\mathsf{uf\text{-}cma}}(\eta, tb+\delta_s, q)+2\cdot n/\eta\cdot q_r\cdot 2^{-\eta}+2\epsilon_P^{\mathsf{prg}}(\eta, tb+\delta_p, q')$ and S, P are respectively the signature scheme and pseudorandom-number generator for protocol $\mathcal{Q}$. The quantities δ_s and δ_p are the respective runtime overheads for the proofs of the **VER** and **FS2** axioms. These are small polynomials and are described in 5.4. Note that each role uses just one signature action. Hence q in (2) is upper bounded by $n_{init} + n_{resp}$, where n_{init} and n_{resp} are the number of instances of the initiator and responder roles respectively. Each role uses just one random number generation action. Hence, q' is upper bounded by $n_{init} + n_{resp}$.

7 Related Work

We discuss (1) formal techniques for proving concrete security, and (2) formal techniques asymptotic security in computational models.

7.1 Formal Concrete Security

Prior work on formal methods for concrete security proofs is largely based on reasoning about program equivalence. This approach operates at the same level of abstraction as traditional proofs in cryptography, where security is proved via a sequence of games, all of which are equivalent up to a small probability. Techniques to prove equivalence of probabilistic programs have been used to formalize cryptographic reduction proofs. Tools for assisting development of such proofs are CryptoVerif [7], CertiCrypt [8], and EasyCrypt [9]. ZooCrypt [13] contains a formal system for computing the concrete bounds of padding-based encryption primitives. In contrast, relational reasoning in QPCL is not exposed to the users of the proof system, and is pushed down to the soundness proofs of axioms related to cryptography. An interesting direction to pursue would be to mechanize the soundness proofs of QPCL's axioms using such a framework.

7.2 Formal Asymptotic Security

While formal techniques for concrete security proofs of cryptography are a relatively new development, a number of approaches have been proposed for deriving proofs of security in asymptotic models, where it is shown that the probability with which a probabilistic polynomial time adversary can break a scheme is a negligible function of the security parameter. We discuss some of these approaches here.

The closest strand of related work to ours include first-order logics such as the computational extensions of PCL [3, 4] and the logic of Bana and Comon-Lundh [14] (CCSA). QPCL borrows heavily from PCL's syntax and reasoning style, which allows for the specification of precise temporal properties and invariants of programs such as the matching-conversations property proved in this paper. This is an important point of difference from CCSA [14], where the

specification language expresses relations between terms appearing in protocol, such as whether one term can be derived from another. An appealing feature of CCSA is that it is designed to be unconditionally sound, in the sense that all assumptions about primitives are encoded in the logic itself. An interesting direction to pursue would be to adapt their unconditionally sound approach to QPCL, while retaining the richness of our specification language.

Symbolic models have been successful in proving the security of a large number protocols, assuming perfect cryptography. In the so-called Dolev-Yao model [15], an attacker's capabilities are defined by a symbolic abstraction of cryptographic primitives. The security of a protocol in a Dolev-Yao model does not, in general, imply security under standard cryptographic assumptions. To justify reasoning with symbolic models, a research direction initiated by Abadi and Rogaway [16] investigates the conditions under which security in the symbolic model implies security in the computational model. These results (e.g., [17, 18, 19, 20, 21]), referred to as *computational soundness proofs*, typically show that under certain conditions a computational attacker can be translated to a symbolic attacker with high probability. Reasoning about cryptography in a symbolic model has been shown to be helpful for automating proofs of protocols [22, 23]. Computational soundness theorems are very strong general results that prove through complicated reductions that any program secure in a particular symbolic model will be asyptotically secure. We conjecture that it is because of the generality of these results that these reductions have not been shown to yield precise concrete bounds for most cryptographic schemes.

8 Conclusion

In this paper, we present Quantitative Protocol Composition Logic (QPCL), a program logic for reasoning about concrete security bounds of of cryptographic protocols. QPCL supports reasoning about temporal trace properties, Hoare logic-style invariants and postconditions of protocol programs, and first-order belief assertions. The semantics and soundness proofs of QPCL depend on a formal probabilistic programming model that describes the concurrent execution of protocol programs with a computationally bounded adversary, while accounting for exact runtimes and probabilities. As an illustrative example, we use QPCL to prove an authentication property with concrete bounds of a challeng-response protocol.

In future work, we hope to extend QPCL to support reasoning about non-trace properties, such as secrecy based on computational indistinguishability, and to cover additional cryptographic primitives (e.g., symmetric and public key encryption, hash functions, and message authentication codes). Because of the correspondence between QPCL presented here and earlier versions of PCL, for which we have proofs for deployed protocols such as SSL and Kerberos, we believe it should be possible to develop exact security analyses for such protocols.

A Appendix

A.1 Proof System and Soundness Proofs

In this section, we prove the soundness of the axioms that we discussed in the main text.

Lemma 7. *The axiom* $\mathbf{PC}_{\uparrow}$ *is sound.*

Proof. Fix a protocol $\mathcal{Q}$, and suppose that $\models_{\mathcal{Q}} B^{\epsilon}\varphi$. We must show that $\models_{\mathcal{Q}} B^{\epsilon}(\psi[P]_I\varphi)$. Given an adversary A, security parameter η, time bound tb, instance vector $\vec{n}$, and assigment V, let $\mathcal{B} = \mathcal{T}_{Q,A}(\eta, tb, \vec{n})$. Since $\models_{\mathcal{Q}} B^{\epsilon}\varphi$, we must have $p\left([\![\neg\varphi]\!]^V(\mathcal{B})\right) \leq \epsilon(\eta, tb, \vec{n})$.

Let $\mathcal{B}_1 = [\![\neg(\psi[P]_I\varphi)]\!]^V(\mathcal{B}) = \{\mathcal{L} \in \mathcal{B} \mid \exists \mathcal{L}_1, \mathcal{L}_2, V'.(\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V), \mathcal{L}_1, V \models \theta, \text{and } \mathcal{L}_2, V' \not\models \varphi\}$. We can assume without loss of generality that the free variables in φ are not bound by $[P]_I$, since this can be ensured via renaming. We can also show by the definition of matching that $V(x) = V'(x)$ for all variables x that are not bound by $[P]_I$. Therefore, we have that $\mathcal{B}_1 = \{\mathcal{L} \in \mathcal{B} \mid \exists \mathcal{L}_1, \mathcal{L}_2, V'.(\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V), \mathcal{L}_1, V \models \theta, \text{and } \mathcal{L}_2, V \not\models \varphi\}$. Additionally, we know using Lemma 1 that φ is a safety property. Also, by the definition of matching, we know that if $(\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V)$, then $\mathcal{L}_2 \subseteq \mathcal{L}$, and consequently, $\mathcal{L}_2, V \not\models \varphi$ implies $\mathcal{L}, V \not\models \varphi$. Therefore, we can show that:

$$
\begin{array}{rcl}
p(\mathcal{B}_1) & = & p(\{\mathcal{L} \in \mathcal{B} \mid \exists \mathcal{L}_1, \mathcal{L}_2, V'.(\mathcal{L} \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V'), \mathcal{L}_1, V \models \theta, \text{ and } \mathcal{L}_2, V \not\models \varphi\}) \\
& \leq & p(\{\mathcal{L} \in \mathcal{B} \mid \exists \mathcal{L}_1, \mathcal{L}_2, V'.(\mathcal{L} \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V'), \mathcal{L}_1, V \models \theta, \text{ and } \mathcal{L}, V \not\models \varphi\}) \\
& \leq & p(\{\mathcal{L} \in \mathcal{B} \mid \mathcal{L}, V \not\models \varphi\}) \\
& = & p\left([\![\neg\varphi]\!]^V(\mathcal{B})\right) \\
& \leq & \epsilon(\eta, tb, \vec{n}).
\end{array}
$$

Therefore, $p\left([\![\neg(\psi[P]_I\varphi)]\!]^V(\mathcal{B})\right) = p(\mathcal{B}_1) \leq \epsilon(\eta, tb, \vec{n})$, and $\models_{\mathcal{Q}} B^{\epsilon}(\psi[P]_I\varphi)$, as desired. □

Lemma 8. *The axiom* $\mathbf{PC}_{\Rightarrow}$ *is sound.*

Proof. We want to show that $\models_{\mathcal{Q}} (\theta[P]_I(\varphi_1 \Rightarrow \varphi_2)) \wedge (\theta[P]_I\varphi_1)) \Rightarrow (\theta[P]_I\varphi_2)$. Given $A, \eta, tb, \vec{n}, V$, let $\mathcal{B} = \mathcal{T}_{\mathcal{Q},A}(\eta, tb, \vec{n})$. Fix a trace $\mathcal{L} \in \mathcal{B}$. Suppose that $\mathcal{L}, V \models \theta[P]_I(\varphi_1 \Rightarrow \varphi_2)$ and $\mathcal{L}, V \models \theta[P]_I\varphi_1$. We must show that $\mathcal{L}, V \models \theta[P]_I\varphi_2$. Fix $\mathcal{L}_1$, $\mathcal{L}_2$ and V' such that $\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}_2, V'$ and $\mathcal{L}_1, V \models \theta$. We must show that $\mathcal{L}_2, V' \models \varphi_2$. This is immediate since it must be the case that $\mathcal{L}_2, V' \models \varphi_1 \Rightarrow \varphi_2$ and $\mathcal{L}_2, V' \models \varphi_1$. □

Lemma 9. *The axiom* $\mathbf{PC}_{\wedge}$ *is sound.*

Proof. This proof is analogous to the previous proof. □

Before proving the soundness of inference rule **HON**, we state three lemmas asserting that the validity of a simple formula φ on a trace is not affected if the trace is extended with a transition in which in an action is not executed, that is, the configuration transition is due to one of the rules abort, wait, or

switch. These can be proved by a straightforward induction on the structure of the formula φ.

Lemma 10. *If trace $\mathcal{L}$ ends in configuration $\mathcal{C}$, $\frac{}{\mathcal{C} \to \mathcal{C}'}(\mathsf{abort})$ and $\mathcal{L}'$ is the extension of $\mathcal{L}$ with $\mathcal{C}'$, then $\mathcal{L}, V \models \varphi$ implies $\mathcal{L}', V \models \varphi$.*

Lemma 11. *If trace $\mathcal{L}$ ends in configuration $\mathcal{C}$, $\frac{}{\mathcal{C} \to \mathcal{C}'}(\mathsf{wait})$ and $\mathcal{L}'$ is the extension of $\mathcal{L}$ with $\mathcal{C}'$, then $\mathcal{L}, V \models \varphi$ implies $\mathcal{L}', V \models \varphi$.*

Lemma 12. *If trace $\mathcal{L}$ ends in configuration $\mathcal{C}$, $\frac{}{\mathcal{C} \to \mathcal{C}'}(\mathsf{switch})$ and $\mathcal{L}'$ is the extension of $\mathcal{L}$ with $\mathcal{C}'$, then $\mathcal{L}, V \models \varphi$ implies $\mathcal{L}', V \models \varphi$.*

Lemma 13. *The inference rule* **HON** *is sound.*

Proof. Suppose that $\models_{\mathcal{Q}} \mathsf{Start}(I) \Rightarrow \varphi$ and $\models_{\mathcal{Q}} \varphi[P]_I\varphi$ for all thread-identifiers I and all roles $P \in IS(\mathcal{Q})$. We must show that $\models_{\mathcal{Q}} \varphi$. Fix $\mathcal{Q}, \mathcal{A}, \eta, tb, \vec{n}, V$, and let $T = \mathcal{T}_{\mathcal{Q},\mathcal{A}}(\eta, tb, \vec{n})$. Consider a trace $\mathcal{L} \in T$, a valuation V, and a thread-identifier I. We show that $\mathcal{L}, V \models \varphi$ by induction on the number of transitions in $\mathcal{L}$.

For the base case, suppose that $\mathcal{L}$ has no transitions (it is an empty trace or has single configuration $\mathcal{C}$), then $\mathcal{L}, V \models \mathsf{Start}(I)$. Since $\models_{\mathcal{Q}} \mathsf{Start}(I) \Rightarrow \varphi$, it follows that $\mathcal{L}, V \models \varphi$, as desired.

If $\mathcal{L}$ has at least one transition, then there exists some trace $\mathcal{L}'$ ending with a configuration $\mathcal{C}$ such that $\mathcal{L}$ is $\mathcal{L}'$ extended with the transition $\mathcal{C} \to \mathcal{C}'$. By the induction hypothesis, we have that $\mathcal{L}' \models \varphi$. We now have the following cases for $\mathcal{C} \to \mathcal{C}'$:

- (act) In this case, $\mathcal{C} \to \mathcal{C}'$ is an action transition. Since traces are generated by executing the actions of protocol programs in order, it can be shown that $\mathcal{L}, V \gg [P]_I|\mathcal{L}_1, \mathcal{L}, V'$, for some $P \in IS(\mathcal{Q})$, and some $\mathcal{L}_1 \prec \mathcal{L}$. By the induction hypothesis, we know that $\mathcal{L}_1, V \models \varphi$. Also, by the antecedent, we know that $\mathcal{L}, V \models \varphi[P]_I\varphi$. Therefore, $\mathcal{L}, V' \models \varphi$. We can assume without loss of generality that the free variables in φ are not bound by $[P]_I$, since this can be ensured via renaming. We can also show by the definition of matching that $V(x) = V'(x)$ for all variables x that are not bound by $[P]_I$. Therefore, we have that $\mathcal{L}, V \models \varphi$.
- (abort) By Lemma 10, $\mathcal{L}, V \models \varphi$.
- (wait) By Lemma 11, $\mathcal{L}, V \models \varphi$.
- (switch) By Lemma 12, $\mathcal{L}, V \models \varphi$.

□

Additional axioms In the main text, we did not give axioms for Contains, Fresh or equality. We now give these axioms. The proof of soundness for each of these axioms is straightforward, and is left to the reader. Axioms **COMP1**-**COMP7** capture computability, and **NCOMP** captures non-computability. Axioms **FS1**-**FS5** capture freshness of nonces. Axioms **EQ1** and **EQ2** capture equality.

COMP1. $B^0\,(\mathsf{Contains}(m, m))$

COMP2. $B^0\,(\mathsf{Contains}(x, y) \wedge \mathsf{Contains}(y, z) \Rightarrow \mathsf{Contains}(x, z))$

COMP3. $B^0\,(\mathsf{Contains}(n, m_0) \Rightarrow \mathsf{Contains}(n, \langle m_0, m_1\rangle))$

COMP4. $B^0\,(\mathsf{Contains}(n, m_1) \Rightarrow \mathsf{Contains}(n, \langle m_0, m_1\rangle))$

COMP5. $B^0\,(\mathsf{Contains}(x, y)\ [P]_I\ \mathsf{Contains}(x, y))$

COMP6. $B^0\,(\mathit{true}\ [m' \leftarrow \mathtt{sign}\ m, k]_I\ \mathsf{Contains}(m, m'))$

COMP7. $B^0\,(\mathit{true}\ [_ \leftarrow \mathtt{verify}\ m', m, k]_I\ \mathsf{Contains}(m', m))$

NCOMP. If m and m' are atomic and $m \neq m'$, then $B^0\,(\neg\mathsf{Contains}(m, m'))$.

FS1. $B^0\,(\mathit{true}\ [n \leftarrow \mathtt{new}]_I\ \mathsf{Fresh}(I, n))$

FS3. $B^0\,(\mathsf{Fresh}(I, n)\ [\mathtt{nosendactions}]_I\mathsf{Fresh}(I, n))$

FS4. $B^0\,(\mathsf{Fresh}(I, n)\ [\mathtt{send}\ m]_I\neg\mathsf{Contains}(n, m) \Rightarrow \mathsf{Fresh}(I, n))$

FS5. $B^0\,(\mathsf{Fresh}(I, n)\ [\mathtt{send}\ m]_I\mathsf{Contains}(n, m) \Rightarrow \mathsf{FirstSend}(I, n, m))$

EQ1. $\langle t_1, \cdots, t_k\rangle = \langle t'_1, \cdots, t'_k\rangle \Rightarrow (t_1 = t'_1) \wedge \cdots \wedge (t_k = t'_k)$

EQ2. $\varphi \wedge (x = t) \Rightarrow [t/x]\varphi$

Our proofs also use axioms for local reasoning about the temporal ordering of actions in a program. In these axioms, $a, a_1, a_2 \in \{\mathtt{send}, \mathtt{receive}, \mathtt{new}, \mathtt{verify}\}$.

AA1. $B^0\,(\mathit{true}[\cdots x \leftarrow a\ t; \cdots]_I\mathsf{a}(I, x, t))$

AA2. $B^0\,(\mathit{true}[\cdots x_1 \leftarrow a_1\ t_1; \cdots x_2 \leftarrow a_2\ t_2 \cdots]_I(\mathsf{a}_1(I, x_1, t_1) < \mathsf{a}_2(I, x_2, t_2)))$

AA3. $\mathsf{a}_1(I, x_1, t_1) < \mathsf{a}_2(I, x_2, t_2) \Rightarrow \mathsf{a}_1(I, x_1, t_1)$

AA4. $\mathsf{a}_1(I, x_1, t_1) < \mathsf{a}_2(I, x_2, t_2) \Rightarrow \mathsf{a}_2(I, x_2, t_2)$

AA5. If for all V, $V(\mathsf{a}(I, x, t)) \neq V(\mathsf{a}'(I, x', t'))$, then $\neg\mathsf{a}(I, x, t)[x' \leftarrow a'\ t']_I\neg\mathsf{a}(I, x, t)$

AA6. $\neg\mathsf{a}(I, x, t)[x' \leftarrow a\ t']_X(\mathsf{a}(I, x, t) \Rightarrow (x' = x) \wedge (t' = t))$

Start. $\mathsf{Start}(I) \Rightarrow \neg\mathsf{a}(I, x, t)$

Informally, axiom **AA1** states that that if a is an action and a is the corresponding action predicate, then after the thread with identifier I executes $(x \leftarrow a\ t;)$, $\mathsf{a}(I, x, t)$ holds. Axiom **AA2** establishes an order between actions in a protocol program. Axioms **AA3** and **AA4** state that if we can establish an order between two actions, both actions must have actually occurred. Axiom **AA5** states that if an action predicate does not unify with an action, and if the action predicate is false before the action, then it is false after the action. Axiom **Start** states that no actions with thread identifier I have occurred if $\mathsf{Start}(I)$ holds on a trace.

A.2 Full Proof of Protocol Example

We now provide a full, formal proof of the correctness of the protocol example.

Using **AA1**,

$$\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\, \mathsf{Verify}(A, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle, \hat{B})$$

is provable. Thus, by **B1** and **B2** so is

$$B^0(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\, \mathsf{Verify}(A, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle, \hat{B})).$$

Choose Δ to be $\mathsf{Honest}(\hat{B})$. By **VER** instantiated to principal $\hat{B}$,

$$\Delta \vdash_{\mathcal{Q}} B^{\epsilon}_{\mathbf{VER}}(\mathsf{Verify}(A, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle, \hat{B}) \Rightarrow \exists \iota.\mathsf{Sign}(\langle \hat{B}, \iota\rangle, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle))$$

is provable. By $\mathbf{PC}_{\uparrow}$, so is

$$\Delta \vdash_{\mathcal{Q}} B^{\epsilon}_{\mathbf{VER}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A(\mathsf{Verify}(A, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle, \hat{B}) \Rightarrow \exists \iota.\mathsf{Sign}(\langle \hat{B}, \iota\rangle, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle))).$$

By $\mathbf{PC}_{\Rightarrow}$, so is

$$B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\ \exists \iota.\mathsf{Sign}(\langle \hat{B}, \iota\rangle, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle)).$$

Let B abbreviate the term $\langle \hat{B}, \iota\rangle$. We have that

$$B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A \exists \iota.\ \mathsf{Sign}(B, s, \langle \text{“}\mathtt{Resp}\text{”}, y, m, \hat{A}\rangle)). \tag{6}$$

Next, we wish to use the **HON** rule to prove the following invariant φ_{inv} about protocol programs.

$$\varphi_{inv} = \mathsf{Sign}(Y, s', \langle \text{“}\mathtt{Resp}\text{”}, y', m', \hat{X}\rangle) \Rightarrow \mathsf{Receive}(Y, m') \wedge \mathsf{FirstSend}(Y, y', \langle y', s'\rangle) \wedge \Phi')$$

where

$$\Phi' \equiv \mathsf{Send}(Y, \langle y', s'\rangle) \Rightarrow \mathsf{Receive}(Y, m') < \mathsf{Send}(Y, \langle y', s'\rangle).$$

Therefore, we need to show that $\mathsf{Start}(X) \Rightarrow \varphi_{inv}$, and for all initial segments P of protocol roles $\varphi_{inv}[P]_X\varphi_{inv}$. By axiom **Start**, we have:

$$\mathsf{Start}(Y) \Rightarrow \neg\mathsf{Sign}(Y, s', \langle \text{“}\mathtt{Resp}\text{”}, y', m', \hat{X}\rangle). \tag{7}$$

And therefore, by straightforward first order reasoning,

$$\mathsf{Start}(Y) \Rightarrow \varphi_{inv}. \tag{8}$$

We need to show for every initial segment P of the roles of the protocol, $\mathsf{Start}(Y)[P]_Y\varphi_{inv}$. We show this for three interesting cases.

We first consider the initial segment $P_1 = \langle x, \hat{X}\rangle \leftarrow \texttt{receive}$. Using **AA5**, we have:

$$\neg\mathsf{Sign}(Y, r, \langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle)[P_1]_Y\neg\mathsf{Sign}(Y, r, \langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle). \tag{9}$$

By straightforward first order reasoning and using **Start** and **G3**, we have:

$$\mathsf{Start}(Y)[P_1]_Y\varphi_{inv}. \tag{10}$$

Next, we consider the initial segment $P_2 = \langle x, \hat{X}\rangle \leftarrow \texttt{receive}; n \leftarrow \texttt{new}; r \leftarrow \texttt{sign}\langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle$. Using **AA1**, we have:

$$true[P_2]_Y\mathsf{Receive}(Y, x). \tag{11}$$

Weakening using **G3**, we have:

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Receive}(Y, x). \tag{12}$$

Also, using **AA5**, we have:

$$\neg\mathsf{Send}(Y, t')[P_2]_Y\neg\mathsf{Send}(Y, t'). \tag{13}$$

Also, by the definition of $\mathsf{FirstSend}$, we can show that:

$$\neg\mathsf{Send}(Y, t') \Rightarrow \mathsf{FirstSend}(Y, n, \langle n, r\rangle). \tag{14}$$

Therefore, using **Start**, **G3** and (13) we have:

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{FirstSend}(Y, n, \langle n, r\rangle). \tag{15}$$

Similarly, we can show that:

$$\mathsf{Start}(Y)[P_2]_Y\neg\mathsf{Send}(Y, \langle r, n\rangle). \tag{16}$$

And therefore, by **G3**,

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Send}(Y, \langle r, n\rangle) \Rightarrow \mathsf{Receive}(Y, x) < \mathsf{Send}(Y, \langle r, n\rangle). \tag{17}$$

Therefore, from (13), (15), (17), and using **G1**, we have:

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Receive}(Y, x)\wedge\mathsf{FirstSend}(Y, n, \langle n, r\rangle)\wedge\mathsf{Send}(Y, \langle r, n\rangle) \Rightarrow \mathsf{Receive}(Y, x) < \mathsf{Send}(Y, \langle r, n\rangle). \tag{18}$$

Now, using **AA5**, **AA6**, and **Start**, for fresh u we can show that:

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Sign}(Y,u) \rightarrow u = \langle r, \langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle\rangle \tag{19}$$

Substituting u with $\langle s', \langle \text{“}\texttt{Resp}\text{”}, y', m', \hat{X}\rangle\rangle$, and using **EQ1**, we have:

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Sign}(Y, \langle s', \langle \text{“}\texttt{Resp}\text{”}, y', m', \hat{X}\rangle\rangle) \Rightarrow (s' = r) \wedge (y' = n) \wedge (m' = x). \tag{20}$$

Applying **G1** to 18 and 20, and rearranging terms, we have:

$$\begin{aligned}&\mathsf{Start}(Y)[P_2]_Y\mathsf{Sign}(Y, \langle s', \langle \text{“}\texttt{Resp}\text{”}, y', m', \hat{X}\rangle\rangle) \Rightarrow \\ &\quad (s' = r) \wedge (y' = n) \wedge (m' = x) \\ &\quad \mathsf{Receive}(Y,x) \wedge \mathsf{FirstSend}(Y, n, \langle n, r\rangle) \wedge \mathsf{Send}(Y, \langle r, n\rangle) \Rightarrow \mathsf{Receive}(Y,x) < \mathsf{Send}(Y, \langle r, n\rangle).\end{aligned} \tag{21}$$

Using **EQ2** on (21), we have what we wanted to prove.

$$\begin{aligned}&\mathsf{Start}(Y)[P_2]_Y\mathsf{Sign}(Y, \langle s', \langle \text{“}\texttt{Resp}\text{”}, y', m', \hat{X}\rangle\rangle) \Rightarrow \\ &\quad \mathsf{Receive}(Y,x) \wedge \mathsf{FirstSend}(Y, s', \langle y', m'\rangle) \wedge \mathsf{Send}(Y, \langle y', s'\rangle) \Rightarrow \mathsf{Receive}(Y,m') < \mathsf{Send}(Y, \langle y', s'\rangle).\end{aligned} \tag{22}$$

Finally, we consider the initial segment $P_3 = \langle x, \hat{X}\rangle \leftarrow \texttt{receive}; n \leftarrow \texttt{new}; r \leftarrow \texttt{sign}\langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle; \texttt{send}\langle \hat{X}, < n, r >\rangle$.

Similar to (12), we can show that

$$\mathsf{Start}(Y)[P_3]_Y\mathsf{Receive}(Y,x). \tag{23}$$

Using **FS1**, we have :

$$true[n \leftarrow \texttt{new}]_Y\mathsf{Fresh}(Y,n). \tag{24}$$

We can also show using **G3** that:

$$\mathsf{Start}(Y)[\langle x, \hat{X}\rangle \leftarrow \texttt{new}]_Y true. \tag{25}$$

Therefore, by **S1**, (24), and (25):

$$\mathsf{Start}(Y)[\langle x, \hat{X}\rangle \leftarrow \texttt{receive}; n \leftarrow \texttt{new}]_Y\mathsf{Fresh}(Y,n). \tag{26}$$

By **FS3**, we have:

$$\mathsf{Fresh}(I,n)[r \leftarrow \texttt{sign}\langle \text{“}\texttt{Resp}\text{”}, n, x, \hat{X}\rangle]\mathsf{Fresh}(I,n) \tag{27}$$

Therefore, by **S1**, (26), and (27):

$$\mathsf{Start}(Y)[P_2]_Y\mathsf{Fresh}(Y,n). \tag{28}$$

Also, using **COMP1** and **COMP3**, we have:

$$\mathsf{Contains}(n, \langle n, r\rangle). \tag{29}$$

Therefore, using **FS5** and (29), we have:

$$\mathsf{Fresh}(Y,n)[\mathtt{send}\langle n,r\rangle)]_Y\mathsf{FirstSend}(\mathsf{Y},\mathsf{n},\langle\mathsf{n},\mathsf{r}\rangle). \tag{30}$$

Therefore, by **S1**, (28), and (30):

$$\mathsf{Start}(Y)[P_3]_Y\mathsf{FirstSend}(Y,n,\langle n,r\rangle). \tag{31}$$

Also, by **AA2**, we have:

$$\mathsf{Start}(Y)[P_3]_Y\mathsf{Receive}(Y,x) < \mathsf{Send}(Y,\langle r,n\rangle) \tag{32}$$

Therefore, we have

$$\mathsf{Start}(Y)[P_3]_Y\mathsf{Send}(Y,\langle r,n\rangle) \Rightarrow \mathsf{Receive}(Y,x) < \mathsf{Send}(Y,\langle r,n\rangle) \tag{33}$$

Therefore, by **G1**, (23), (31), and (33).

$$\mathsf{Start}(Y)[P_3]_Y\mathsf{Receive}(Y,x)\wedge\mathsf{FirstSend}(Y,n,\langle n,r\rangle)\wedge\mathsf{Send}(Y,\langle r,n\rangle) \Rightarrow \mathsf{Receive}(Y,x) < \mathsf{Send}(Y,\langle r,n\rangle). \tag{34}$$

Analogously, to proof steps of (18)-(21), we can use 34 to show:

$$\begin{aligned}&\mathsf{Start}(Y)[P_3]_Y\mathsf{Sign}(Y,\langle s',\langle\text{“}\mathtt{Resp}\text{”},y',m',\hat{X}\rangle\rangle) \Rightarrow\\ &\quad \mathsf{Receive}(Y,x)\wedge\mathsf{FirstSend}(Y,s',\langle y',m'\rangle)\wedge\mathsf{Send}(Y,\langle y',s'\rangle) \Rightarrow \mathsf{Receive}(Y,m') < \mathsf{Send}(Y,\langle y',s'\rangle).\end{aligned} \tag{35}$$

We can show $\mathsf{Start}(Y)[P]_Y\varphi_{inv}$ for every initial sequence, and by using **HON**, we obtain the following invariant about protocols:

$$B^0(\mathsf{Sign}(Y,s',\langle\text{“}\mathtt{Resp}\text{”},y',m',\hat{X}\rangle) \Rightarrow \mathsf{Receive}(Y,m')\wedge\mathsf{FirstSend}(Y,y',\langle y',s'\rangle)\wedge\Phi'). \tag{36}$$

By instantiating some of the free variables in (36) and using axiom B1, we get that

$$B^0(\mathsf{Sign}(B,s,\langle\text{“}\mathtt{Resp}\text{”},y,m,\hat{A}\rangle) \Rightarrow \mathsf{Receive}(B,m)\wedge\mathsf{FirstSend}(B,y,\langle y,r\rangle)\wedge\Phi)$$

is provable, where

$$\Phi \equiv \mathsf{Send}(B,\langle y,s\rangle) \Rightarrow \mathsf{Receive}(B,m) < \mathsf{Send}(B,\langle y,s\rangle).$$

By $\mathbf{PC}_{\uparrow}$,

$$B^0(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\ \mathsf{Sign}(B,s,\langle\text{“}\mathtt{Resp}\text{”},y,m,\hat{A}\rangle) \Rightarrow (\mathsf{Receive}(B,m)\wedge\mathsf{FirstSend}(B,y,\langle y,s\rangle)\wedge\Phi))$$

is provable. By $\mathbf{PC}_{\Rightarrow}$ and (6),

$$\Delta \vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\ \exists\iota.\ (\mathsf{Receive}(B,m)\wedge\mathsf{FirstSend}(B,y,\langle y,s\rangle)\wedge\Phi))$$

is provable. By **FS1**, **COMP** so is

$$B^0(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A(\mathsf{FirstSend}(A,m,m)\wedge\mathsf{Contains}(m,m)\wedge\mathsf{Contains}(y,\langle y,s\rangle))).$$

Therefore by $\mathbf{PC}_\wedge$ and B1, so is

$$\Delta\vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\exists\iota.\ (\mathsf{Receive}(B,m)\wedge\mathsf{FirstSend}(A,m,m)\wedge\mathsf{Contains}(m,m)\wedge \mathsf{Receive}(A,\langle y,s\rangle)\wedge\mathsf{FirstSend}(B,y,\langle y,s\rangle)\wedge\mathsf{Contains}(y,\langle y,s\rangle)\wedge\Phi)). \tag{37}$$

We instantiate **FS2** with message m and nonce m, and use this to establish $\mathsf{Send}(A,m) < \mathsf{Receive}(B,m)$:

$$B^{\epsilon_{\mathbf{FS2}}}(\mathsf{FirstSend}(A,m,m)\wedge\mathsf{Receive}(B,m)\wedge\mathsf{Contains}(m,m)\Rightarrow\mathsf{Send}(A,m)<\mathsf{Receive}(B,m)) \tag{38}$$

Similarly, we instantiate **FS2** with message $\langle y,s\rangle$ and nonce s. We will use this to establish $\mathsf{Send}(B,\langle y,s\rangle) < \mathsf{Receive}(A,\langle y,s\rangle)$.

$$B^{\epsilon_{\mathbf{FS2}}}(\mathsf{FirstSend}(B,y,\langle y,s\rangle)\wedge\mathsf{Receive}(B,\langle y,s\rangle)\wedge\mathsf{Contains}(y,\langle y,s\rangle)\Rightarrow\mathsf{Send}(B,\langle y,s\rangle)<\mathsf{Receive}(A,\langle y,s\rangle)). \tag{39}$$

Therefore, by the conjunction of (38) and (39) and using **B3**:

$$B^{2\cdot\epsilon_{\mathbf{FS2}}}(\mathsf{FirstSend}(A,m,m)\wedge\mathsf{Receive}(B,m)\wedge\mathsf{Contains}(m,m)\wedge\mathsf{FirstSend}(B,y,\langle y,s\rangle) \wedge(\mathsf{Receive}(B,\langle y,s\rangle)\wedge\mathsf{Contains}(y,\langle y,s\rangle) \Rightarrow\mathsf{Send}(A,m)<\mathsf{Receive}(B,m))\wedge\mathsf{Send}(B,\langle y,s\rangle)<\mathsf{Receive}(A,\langle y,s\rangle))). \tag{40}$$

Therefore by (37), (40), and using $\mathbf{PC}_\uparrow$ and $\mathbf{PC}_\Rightarrow$,

$$\Delta\vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}+2\cdot\epsilon_{\mathbf{FS2}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\exists\iota.\ \mathsf{Send}(A,m)<\mathsf{Receive}(B,m)\wedge\mathsf{Send}(B,\langle y,s\rangle)<\mathsf{Receive}(A,\langle y,s\rangle)\wedge\Phi))$$

Finally, by **AA4** and $\mathbf{PC}_\Rightarrow$ on Φ,

$$\Delta\vdash_{\mathcal{Q}} B^{\epsilon_{\mathbf{VER}}+2\cdot\epsilon_{\mathbf{FS2}}}(\mathit{true}[\mathbf{Init}_{CR}(\hat{B})]_A\exists\iota.\ (\mathsf{Send}(A,m)<\mathsf{Receive}(B,m) \wedge\mathsf{Receive}(B,m)<\mathsf{Send}(B,\langle y,s\rangle)\wedge\mathsf{Send}(B,\langle y,s\rangle)<\mathsf{Receive}(A,\langle y,s\rangle))).$$

A.3 A First-Order Logic of Belief

Halpern [1] presents a "quantitative" variant of first-order conditional logic that includes formulas whose semantics capture the form of conditional probability statements that are common in cryptographic security definitions. In this paper, we use a simplified version of that logic as a foundation, and show how we can devise an elegant, powerful variant of the logic considered in [3]. We summarize below the relevant results of [1].

The syntax of (qualitative) first-order conditional logic is straightforward. Fix a finite first-order vocabulary $\mathcal{V}$ consisting, as usual, of function symbols, predicate symbols, and constants. Starting with atomic formulas of first-order logic over the vocabulary $\mathcal{V}$, we form more complicated formulas by closing

off under the standard truth-functional connectives (i.e., $\wedge$,$\vee$, $\neg$, and $\Rightarrow$), first-order quantification, and the binary modal operator $\rightarrow$. Thus, a typical formula is $\forall x(P(x) \rightarrow \exists y(Q(x,y) \rightarrow R(y)))$. Let $\mathcal{L}_C(\mathcal{V})$ be the resulting language. We are most interested in the fragment $\mathcal{L}_C^0(\mathcal{V})$ of $\mathcal{L}_C(\mathcal{V})$ consisting of all formulas in $\mathcal{L}_C(\mathcal{V})$ that have the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$, where φ and ψ are first-order formulas. (We henceforth omit the $\mathcal{V}$ unless it is necessary for clarity.) Note that in $\mathcal{L}_C^0$ we cannot negate an $\rightarrow$ formula, nor can we take the conjunction of two $\rightarrow$ formulas. As we shall see, the fact that we cannot take the conjunction is not a real lack of expressive power, since we consider sets of formulas (where a set can be identified with the conjunction of the formulas in the set). However, the lack of negation does make for a loss of expressive power.

There are many approaches to giving semantics to first-order conditional logic (see [24] for an overview). The one most relevant here is what has been called ϵ-semantics [2]. Under this approach, the semantics of formulas in $\mathcal{L}_C(\mathcal{V})$ is given with respect to *PS structures*. A PS structure is a tuple $M = (D, W, \pi, \mathcal{P})$, where D is a domain, W is a set of worlds, π is an *interpretation* that associates with each predicate symbol (resp., function symbol, constant symbol) in $\mathcal{V}$ and world $w \in W$ a predicate (resp., function, domain element) of the right arity, and $\mathcal{P} = \langle \Pr_1, \Pr_2, \ldots \rangle$ is a probability sequence. As usual, a *valuation* V associates with each variable x an element $V(x) \in D$.

Given a valuation V, we give semantics to $\wedge$, $\neg$, $\Rightarrow$, and $\forall$ in the standard way. In particular, we determine whether a first-order formula φ is true at a world $w \in W$ as usual. In this case, we write $(M, V, w) \models^{\mathsf{P}} \varphi$. Let $[\![\varphi]\!]_{M,V} = \{w : (M, V, w) \models^{\mathsf{P}} \varphi\}$. If φ is a closed formula, so that its truth does not depend on the valuation, we occasionally write $[\![\varphi]\!]_M$ rather than $[\![\varphi]\!]_{M,V}$. We write $(M, V) \models^{\mathsf{P}} \varphi$ if $(M, V, w) \models^{\mathsf{P}} \varphi$ for all worlds w.

The truth of an $\rightarrow$ formula does not depend on the world, but only on the PS structure. We can consider two possible semantics for conditional formulas in PS structures. The first just asks for convergence:

$$(M, V) \models^{\mathsf{P}} \varphi \rightarrow \psi \text{ if } \lim_{n \rightarrow \infty} \Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) = 1,$$

where $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V})$ is taken to be 1 if $\Pr_n([\![\varphi]\!]_{M,V}) = 0$.

The second approach may be of more interest in cryptographic applications, and asks for super-polynomial convergence. In this approach, we define

> $(M, V) \models^{\mathsf{P}} \varphi \rightarrow \psi$ if for all $k \geq 1$, there exists some $n_k \geq 0$ such that, for all $n \geq n_k$, $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - (1/n^k)$.[11]

It turns out that the two appraoches are characterized by the same axiom system, so we do not distinguish them further here. As usual, we write $M \models^{\mathsf{P}} \varphi$ if $(M, V) \models^{\mathsf{P}} \varphi$ for all valuations V, and $\mathcal{M} \models^{\mathsf{P}} \varphi$ if $M \models^{\mathsf{P}} \varphi$ for all PS structures in a set $\mathcal{M}$.

[11] Note that it follows that if $p(n)$ is a polynomial whose leading coefficient is positive, then $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - (1/p(n))$ for sufficiently large n. In [1], this was the condition used to define the semantics of $\varphi \rightarrow \psi$. The approach used here is simpler, but equivalent.

In this paper, we need only formulas where the antecedent of $\rightarrow$ is *true*. We take $B\varphi$ (which can be read "φ is believed" or "the agent believes φ') to be an abbreviation of $true \rightarrow \varphi$. Roughly speaking, $B\varphi$ holds if the probability of φ approaches 1 faster than any inverse polynomial. Let $\mathcal{L}_C^B$ be the fragment of $\mathcal{L}_C^0$ where $\rightarrow$ occurs only in the context B. As we have seen, such formulas are used in **VER**; they are also needed in, for example, the CMA security definition for digital signatures [25]).

Sound and complete axiomatizations for first-order conditional logic are well known [26]. Halpern [1] gives a sound and complete axiomatization for $\mathcal{L}_C^0$. The advantage of considering this fragment is that it lends itself naturally to quantitative reasoning. We now review the axiomatization of $\mathcal{L}_C^0$, show how to specialize it to $\mathcal{L}_C^B$, and give the quantitative version of the axioms.

A.4 A Sound and Complete Axiomatization for First-Order Conditional Logic

To describe the axioms, we define two more languages. Let $\mathcal{L}^{fo}$ be the pure first-order fragment of $\mathcal{L}_C$, consisting of $\rightarrow$-free formulas; let $\mathcal{L}_C^-$ be the fragment of $\mathcal{L}_C^0$ consisting of all formulas of the form $\varphi \rightarrow \psi$ where φ and ψ are closed first-order formulas. The axiomatization is based on system $\mathbf{P}_\Lambda$, also known as *the KLM properties* [27], restricted to models where all worlds satisfy a particular first-order theory Λ. The system $\mathbf{P}_\Lambda$ applies only to formulas in $\mathcal{L}_C^-$ (so $\rightarrow$ formulas cannot be universally quantified); we will need additional axioms to deal with quantification.

We state the axioms in terms of statements of the form $\Delta \hookrightarrow \varphi$, where $\Delta \subseteq \mathcal{L}^{fo} \cup \mathcal{L}_C^0$ and $\varphi \in \mathcal{L}_C^0$. Roughly speaking, $\Delta \hookrightarrow \varphi$ is interpreted as "φ follows from the formulas in Δ". We write $(M, V) \models^{\mathsf{p}} \Delta \hookrightarrow \varphi$ if $(M, V) \models^{\mathsf{p}} \varphi'$ for every formula $\varphi' \in \Delta$ implies that $(M, V) \models^{\mathsf{p}} \varphi$. A collection Δ of formulas in $\mathcal{L}_C^0$ can be written as $\Delta_\rightarrow \cup \Delta_{fo}$, where $\Delta_\rightarrow \subseteq \mathcal{L}_C^0$ and $\Delta_{fo} \subseteq \mathcal{L}^{fo}$.

We first describe the axioms of $\mathbf{P}_\Lambda$, adapted to this framework. The axioms all have conclusions of the form $\varphi \rightarrow \psi$, since they are intended for reasoning about $\mathcal{L}_C^-$. Let $\vdash_\Lambda$ denote provability in first-order logic given the axioms in the theory Λ. (In describing these axioms, we keep the naming convention used in [1], which in turn is based on [27], for consistency.)

LLE$^+$. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi_1 \Leftrightarrow \varphi_2$, then from $\Delta \hookrightarrow \varphi_1 \rightarrow \psi$ infer $\Delta \hookrightarrow \varphi_2 \rightarrow \psi$ (left logical equivalence).

RW$^+$. If $\vdash_{\Lambda \cup \Delta_{fo}} \psi_1 \Rightarrow \psi_2$, then from $\Delta \hookrightarrow \varphi \rightarrow \psi_1$ infer $\Delta \hookrightarrow \varphi \rightarrow \psi_2$ (right weakening).

REF. $\Delta \hookrightarrow \varphi \rightarrow \varphi$ (reflexivity).

AND. From $\Delta \hookrightarrow \varphi \rightarrow \psi_1$ and $\Delta \hookrightarrow \varphi \rightarrow \psi_2$ infer $\Delta \hookrightarrow \varphi \rightarrow \psi_1 \wedge \psi_2$.

OR. From $\Delta \hookrightarrow \varphi_1 \rightarrow \psi$ and $\Delta \hookrightarrow \varphi_2 \rightarrow \psi$ infer $\Delta \hookrightarrow \varphi_1 \vee \varphi_2 \rightarrow \psi$.

CM. From $\Delta \hookrightarrow \varphi_1 \rightarrow \varphi_2$ and $\Delta \hookrightarrow \varphi_1 \rightarrow \psi$ infer $\Delta \hookrightarrow \varphi_1 \wedge \varphi_2 \rightarrow \psi$ (cautious monotonicity).

It is interesting to contrast LLE$^+$, RW$^+$, and CM. While RW$^+$ allows a formula ψ_1 on the right-hand side of $\rightarrow$ to be replaced by a weaker formula ψ_2 (that is, a formula such that that $\psi_1 \Rightarrow \psi_2$ is provable), LLE$^+$ just allows a formula φ_1 on the left-hand side of $\rightarrow$ to be replaced by an equivalent formula φ_2, rather than a stronger formula. CM allows the replacement of a formula φ_1 on the left-hand side by a stronger formula, $\varphi_1 \wedge \varphi_2$, but only if $\varphi_1 \rightarrow \varphi_2$ holds. Intuitively, this says that if φ_2 and ψ each almost always hold given φ_1, then ψ almost always holds given both φ_1 and φ_2. Monotonicity does not hold in general. That is, if $\varphi_1 \Rightarrow \varphi_2$ is provable and $\varphi_2 \rightarrow \psi$ holds, then $\varphi_1 \rightarrow \psi$ does not necessarily hold. For a simple counterexample, it is not the case that if $\mathit{true} \rightarrow \psi$ holds then $\neg\psi \rightarrow \psi$ holds. If $\psi(x)$ states that x cannot forge a signature, we might expect that almost always, x cannot forge the signature ($\mathit{true} \rightarrow \psi(x)$), but it surely is not the case that x cannot forge the signature given that x can forge it. The other two axioms AND and OR are easy to explain. AND says that if both ψ_1 and ψ_2 almost always hold given φ, then so does $\psi_1 \wedge \psi_2$, while OR allows reasoning by cases: if ψ almost always holds given each of φ_1 and φ_2, then it almost always holds given their disjunction.

The following additional axioms account for universal quantification, and the fact that $\mathcal{L}^0_C$ includes first-order formulas. One of the axiom uses the notion of *interpretation-independent* formula, which is a first-order formula that does not mention any constant, function, or predicate symbols (and thus, is a formula whose atomic predicates are all of the form $x = y$). In the axioms below, we assume that Λ is a first-order theory.

Λ-AX$^+$. If $\varphi \in \mathcal{L}^{fo}$ and $\vdash_{\Lambda \cup \Delta_{fo}} \varphi$, then from $(\Delta \cup \{\varphi\}) \hookrightarrow \psi$ infer $\Delta \hookrightarrow \psi$.[12]

F1$^+$. If z is a variable that does not appear in φ, then from $(\Delta \cup \varphi[x/z]) \hookrightarrow \psi$ infer $(\Delta \cup \forall x\varphi) \hookrightarrow \psi$.

F3$^+$. If x does not appear free in Δ, then from $\Delta \hookrightarrow \varphi$ infer $\Delta \hookrightarrow \forall x\varphi$.

EQ. If x does not appear free in Δ, φ, or ψ, and σ is a first-order formula, then from $(\Delta \cup \{\sigma\}) \hookrightarrow \varphi$ infer $(\Delta \cup \{\exists x\sigma\}) \hookrightarrow \varphi$ (existential quantification).

REN. If $y_1, \ldots, y_n$ do not appear in φ, then from $\Delta \hookrightarrow \forall x_1 \ldots \forall x_n\varphi$ infer $\Delta \hookrightarrow \forall y_1 \ldots \forall y_n(\varphi[x_1/y_1, \ldots, x_n/y_n])$ (renaming).

II. If σ_1 and σ_2 are interpretation-independent, then from $(\Delta \cup \{\sigma_1\}) \hookrightarrow \varphi$ and $(\Delta \cup \{\sigma_2\}) \hookrightarrow \varphi$, infer $(\Delta \cup \{\sigma_1 \vee \sigma_2\}) \hookrightarrow \varphi$ (interpretation independence).

Let $\mathbf{P}^+_\Lambda$ be the axiom system consisting of LLE$^+$, RW$^+$, REF, AND, OR, CM, Λ-AX$^+$, F1$^+$, F3$^+$, EQ, REN, and II. A *derivation from* $\mathbf{P}^+_\Lambda$ consists of a sequence of steps, where each step has the form $\Delta \hookrightarrow \varphi$ and either (a) $\varphi \in \Delta$ or (b) there is an inference rule "if condition C holds, then from $\Delta_1 \hookrightarrow \varphi_1$, ...,

[12] This rule has changed somewhat from the one in [1], to make it more convenient to apply.

$\Delta_k \hookrightarrow \varphi_k$ infer $\Delta \hookrightarrow \varphi$", condition C holds, and $\Delta_1 \hookrightarrow \varphi_1, \ldots, \Delta_k \hookrightarrow \varphi_k$ appear in earlier steps of the derivation. (The "condition C" might be a statement such as "$\vdash_{\Lambda \cup \Delta_{fo}} \varphi \Leftrightarrow \psi$ and $\varphi \in \mathcal{L}^{fo}$", which appears in LLE$^+$.) We write $\mathbf{P}^+_\Lambda \vdash \Delta \hookrightarrow \varphi$ if there is a derivation from $\mathbf{P}^+_\Lambda$ whose last line in $\Delta \hookrightarrow \varphi$.

Let $\mathcal{PS}(\Lambda)$ consist of all PS structures M such that world in M satisfies Λ. We write $\mathcal{PS}(\Lambda) \models^{\mathsf{p}} \Delta \hookrightarrow \varphi$ if $(M, V) \models^{\mathsf{p}} \Delta \hookrightarrow \varphi$ for all PS structures $M \in \mathcal{PS}(\Lambda)$ and valuations V.

Theorem 2. *[1] If* $\Delta \cup \{\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)\} \subseteq \mathcal{L}^0_C$, *then* $\mathbf{P}^B_\Lambda \vdash \Delta \hookrightarrow \forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$ *if and only if* $\mathcal{PS}(\Lambda) \models^{\mathsf{p}} \Delta \hookrightarrow \forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$.

We now specialize these axioms to $\mathcal{L}^B_C$. We start with the analogues of axioms from $\mathbf{P}_\Lambda$. Consider the following three axioms:

B1. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi \Rightarrow \varphi'$, then from $\Delta \hookrightarrow B\varphi$ infer $\Delta \hookrightarrow B\varphi'$.

B2. $\Delta \hookrightarrow B(\mathit{true})$.

B3. From $\Delta \hookrightarrow B\varphi$ and $\Delta \hookrightarrow B\varphi'$ infer $\Delta \hookrightarrow B(\varphi \wedge \varphi')$.

It should be clear that B1 is just RW$^+$ restricted to $\mathcal{L}^B_C$; B2 is REF restricted to the only formula in $\mathcal{L}^B_C$ to which it applies, namely $\mathit{true} \rightarrow \mathit{true}$; and B3 is the restriction of AND to $\mathcal{L}^B_C$. It is easy to see that there are no instances of LLE$^+$, OR, or CM in $\mathcal{L}^B_C$. The version of B1–B3 in the main text can be viewed as the special case where $\Delta = \emptyset$.

Let $\mathbf{P}^B_\Lambda$ be the result of replacing LLE$^+$, RW$^+$, REF, AND, OR, and CM in $\mathbf{P}^+_\Lambda$ by B1–B3. Given a collection Δ of $\rightarrow$ formulas, we write $\mathbf{P}^B_\Lambda \vdash \Delta \hookrightarrow \varphi$ if $\varphi \rightarrow \psi$ can be derived from Δ using these rules. A *derivation from* Δ consists of a sequence of steps of the form $\Delta \hookrightarrow \varphi$, where either (a) $\varphi \in \Delta$, (b) φ is $B_i(\mathit{true})$ (so that this is just axiom B2), or (c) $\Delta \hookrightarrow \varphi$ follows from previous steps by an application of B1 or B3.

Theorem 3. *If* $\Delta \subseteq \mathcal{L}^B_C \cup \mathcal{L}^{fo}$ *and* $\forall x_1 \ldots \forall x_n B\varphi \in \mathcal{L}^B_C$, *then* $\mathbf{P}^B_\Lambda \vdash \Delta \hookrightarrow \forall x_1 \ldots \forall x_n B\varphi$ *if and only if* $\mathcal{PS}(\Lambda) \models^{\mathsf{p}} \Delta \hookrightarrow \forall x_1 \ldots \forall x_n B\varphi$.

Proof. The soundness follows from the soundness of $\mathbf{P}^+_\Lambda$, proved in [1]. To prove completeness, we first prove that B1-B3 are sound and complete for the fragment $\mathcal{L}^B_C \cap \mathcal{L}^-_C$. Again, soundness is immediate. Completeness is also quite straightforward, as we now show.

We start by showing that $\mathcal{PS}(\Lambda) \models^{\mathsf{p}} \{\psi_1, \ldots, \psi_k, B\varphi_1, \ldots, B\varphi_m\} \hookrightarrow B\sigma$ (where $\psi_1, \ldots, \psi_k, \varphi_1, \ldots, \varphi_m$, and σ are all in $\mathcal{L}^{fo}$) iff $\Lambda \models^{\mathsf{p}} (\psi_1 \wedge \ldots \wedge \psi_k \wedge \varphi_1 \wedge \ldots \wedge \varphi_m) \Rightarrow \sigma$. The "if" direction is obvious. For the "only if" direction, suppose that $\Lambda \not\models^{\mathsf{p}} (\psi_1 \wedge \ldots \wedge \psi_k \wedge \varphi_1 \wedge ... \wedge \varphi_m) \Rightarrow \sigma$. Then $\sigma' = \psi_1 \wedge \ldots \wedge \psi_k \wedge \varphi_1 \wedge \ldots \wedge \varphi_m \wedge \neg\sigma$ is satisfied in a relational structure R that also satisfies all the formulas in Λ. Let $M = (D, \{w\}, \pi, \mathcal{P})$, where D is the domain in R, $\pi(w)$ is the interpretation in R, and $\mathcal{P} = (\Pr, \Pr, \Pr, \ldots)$, and Pr assigns probability 1 to world w. Then $M \in \mathcal{PS}(\Lambda)$ and $M \models^{\mathsf{p}} \psi_1 \wedge \ldots \wedge \psi_k \wedge B\varphi_1 \wedge \ldots \wedge B\varphi_m$, but $M \models^{\mathsf{p}} \neg B\sigma$. Thus, $M \not\models^{\mathsf{p}} \{\psi_1, \ldots, \psi_k, B\varphi_1, \ldots, B\varphi_m\} \hookrightarrow B\sigma$.

Now suppose that $\mathcal{PS}(\Lambda) \models^{\mathsf{P}} \Delta \hookrightarrow B\varphi$, where $\Delta = \{\psi_1, \ldots, \psi_k, B\varphi_1, \ldots, B\varphi_m\}$. Since $\vdash_{\Lambda \cup \Delta_{fo}} \mathit{true} \Rightarrow (\psi_1 \wedge \ldots \wedge \psi_k)$, by B1 and B2, it follows that $\mathbf{P}^B_\Lambda \vdash \Delta \hookrightarrow B(\psi_1 \wedge \ldots \wedge \psi_k)$. Since $B\varphi_1, \ldots, B\varphi_m \in \Delta$, applying B3 repeatedly, it follows that $\mathbf{P}^B_\Lambda \vdash B(\psi_1 \wedge \ldots \wedge \psi_k \wedge \varphi_1 \wedge \ldots \wedge \varphi_m)$. By the arguments above and the completeness of first-order logic, since $\mathcal{PS}(\Lambda) \models^{\mathsf{P}} \Delta \hookrightarrow B\varphi$, we must have $\vdash_\Lambda (\psi_1 \wedge \ldots \wedge \psi_k \wedge \varphi_1 \wedge \ldots \wedge \varphi_m) \Rightarrow \sigma$. Thus, applying B1 again, it follows that $\mathbf{P}^B_\Lambda \vdash \Delta \hookrightarrow B\varphi$.

Given the soundness and completeness of B1-B3 for the $\mathcal{L}^-_C \cap \mathcal{L}^B$, the proof that adding the remaining axioms gives soundness and completeness $\mathcal{L}^B$ is now identical to that going from $\mathcal{L}^-_C$ to $\mathcal{L}_C$ given in Theorem 3.5 of [1]. We omit further details here. □

A.5 Quantitative reasoning

Up to now, we have just considered a qualitative semantics for conditional logic. A formula of the form $\varphi \rightarrow \psi$ is true if, for all k, the conditional probability of ψ given φ is at least $1 - 1/n^k$ for all sufficiently large n. Here n can be, for example, the security parameter. While this asymptotic complexity certainly gives insight into the security of a protocol, in practice, a system designer wants to achieve a certain level of security, and needs to know, for example, how large to take the keys in order to achieve this level. In [1], a more quantitative semantics appropriate for such reasoning was introduced and related to the more qualitative "asymptotic" semantics. We briefly review the relevant details here, and specialize them to the language of belief.

The syntax of the quantitative language considered in [1], denoted $\mathcal{L}_{C,q}$, is just like that of $\mathcal{L}^0_C$, except that instead of formulas of the form $\varphi \rightarrow \psi$ there are formulas of the form $\varphi \rightarrow^r \psi$, where r is a real number in $[0, 1]$. The semantics of such a formula is straightforward:

> $(M, V) \models \varphi \rightarrow^r \psi$ if there exists some $n^* \geq 0$ such that
> for all $n \geq n^*$, $\Pr_n(\llbracket\psi\rrbracket_{M,V} \mid \llbracket\varphi\rrbracket_{M,V}) \geq 1 - r$).

In our quantitative logic, we do not have a sequence of probability distributions on the left hand-side of $\models^{\mathsf{P}}$; rather, we have a single distribution (determined by the function ϵ and the parameters η tb, and $\vec{\eta}$). Nevertheless, as we shall see, the axioms that characterize quantitative conditional logic in the case of a sequence of distributions are essentially the same as those for a single distribution.

Let $\mathcal{L}^0_{C,q}$ be the obvious analogue of $\mathcal{L}^0_C$, consisting of all formulas of the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow^r \psi)$, where φ and ψ are first-order formulas. For each of the axioms and rules in system $\mathbf{P}_\Lambda$, there is a corresponding sound axiom or rule in $\mathcal{L}^0_{C,q}$:

LLE^q. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi_1 \Leftrightarrow \varphi_2$, then from $\Delta \hookrightarrow \varphi_1 \rightarrow^r \psi$ infer $\Delta \hookrightarrow \varphi_2 \rightarrow^r \psi$.

RW^q. If $\vdash_{\Lambda \cup \Delta_{fo}} \psi_1 \Rightarrow \psi_2$, then from $\Delta \hookrightarrow \varphi \rightarrow^r \psi_1$ infer $\Delta \hookrightarrow \varphi \rightarrow^r \psi_2$.

REF^q. $\Delta \hookrightarrow \varphi \rightarrow^0 \varphi$ (reflexivity).

AND^q. From $\Delta \hookrightarrow \varphi \rightarrow^{r_1} \psi_1$ and $\Delta \hookrightarrow \varphi \rightarrow^{r_2} \psi_2$ infer $\Delta \hookrightarrow \varphi \rightarrow^{r_3} \psi_1 \wedge \psi_2$, where $r_3 = \min(r_1 + r_2, 1)$.

OR^q. From $\Delta \hookrightarrow \varphi_1 \rightarrow^{r_1} \psi$ and $\Delta \hookrightarrow \varphi_1 \rightarrow^{r_2} \psi$ infer $\Delta \hookrightarrow \varphi_1 \vee \varphi_2 \rightarrow^{r_3} \psi$, where $r_3 = \min(\max(2r_1, 2r_2), 1)$.

CM^q. From $\Delta \hookrightarrow \varphi_1 \rightarrow^{r_1} \varphi_2$ and $\Delta \hookrightarrow \varphi_1 \rightarrow^{r_2} \psi$ infer $\Delta \hookrightarrow \varphi \wedge \varphi_2 \rightarrow^{r_3} \psi$, where $r_3 = \min(r_1 + r_2, 1)$.

Let $\mathbf{P}^{+,q}_\Lambda$ consist of the rules above, together with F1$^+$, F3$^+$, EQ, REN, and II (all of which hold with no change in the quantitative setting), and

INC. If $r_1 \leq r_2$, then from $\Delta \hookrightarrow \varphi \rightarrow^{r_1} \psi$ infer $\Delta \hookrightarrow \varphi \rightarrow^{r_2} \psi$.

Theorem 4. *[1] The rules in $\mathbf{P}^{+,q}_\Lambda$ are all sound.*

There is a deep relationship between $\mathbf{P}^+_\Lambda$ and $\mathbf{P}^{+,q}_\Lambda$. To make it precise, given a set of formulas $\Delta \subseteq \mathcal{L}^0_C$, say that $\Delta' \subseteq \mathcal{L}^0_{C,q}$ is a *quantitative instantiation* of Δ if, for every formula $\varphi \rightarrow \psi \in \Delta$, there is a bijection f from Δ to Δ' such that, for every formula $\varphi \rightarrow \psi \in \Delta$, there is a real number $r \in [0, 1]$ such that $f(\varphi \rightarrow \psi) = \varphi \rightarrow^r \psi$. That is, Δ' is a quantitative instantiation of Δ if each qualitative formula in Δ has a quantitative analogue in Δ'.

The key result in [1] is the following, which shows the power of using of $\mathbf{P}^+_\Lambda$. Specifically, it shows that if $\Delta \hookrightarrow \varphi \rightarrow \psi$ is derivable in $\mathbf{P}^+_\Lambda$ then, for all $r \in [0, 1]$, there exists a quantitative instantiation Δ' of Δ such that $\varphi \rightarrow^r \psi$ is derivable from Δ' in $\mathbf{P}^{+,q}_\Lambda$. Thus, if the system designer wants security at level r (that is, she wants to know that the desired security property holds with probability at least $1 - r$), then if she has a qualitative proof of the result, she can compute the strength with which her assumptions must hold in order for the desired conclusion to hold. For example, she can compute how to set the security parameters in order to get the desired level of security. This result can be viewed as justifying qualitative reasoning. Roughly speaking, it says that it is safe to avoid thinking about the quantitative details, since they can always be derived later. Note that this result would not hold if the language allowed negation. For example, even if $\neg(\varphi \rightarrow \psi)$ could be proved given some assumptions (using the axiom system AX^Λ_C), it would not necessarily follow that $\neg(\varphi \rightarrow^q \psi)$ holds, even if the probability of the assumptions was taken arbitrarily close to one.

Theorem 5. *[1] If $\mathbf{P}^+_\Lambda \vdash \Delta \hookrightarrow \varphi \rightarrow \psi$, then for all $r \in [0, 1]$, there exists a quantitative instantiation Δ' of Δ such that $\mathbf{P}^{+,q}_\Lambda \vdash \Delta' \hookrightarrow \varphi \rightarrow^q \psi$. Moreover, Δ' can be found in polynomial time, given the derivation of $\Delta \hookrightarrow \varphi \rightarrow \psi$.*

Intuitively, Theorem 5 says that if there is a qualitative proof of $\varphi \rightarrow \psi$ from the assumptions in Δ, given a desired confidence level $1 - r$, we can, in polynomial time, compute the probability with which the assumptions in Δ need to hold in order to ensure that the conclusion $\varphi \rightarrow \psi$ holds with confidence $1-r$.

Let $\mathcal{L}^{B,q}_C$ be the restriction of $\mathcal{L}_{C,q}$ so that $\rightarrow^r$ appears only in the context of *true* $\rightarrow^r \varphi$, which we abbreviate as $B^r\varphi$. The quantitative analogue of B1–B3 is obvious:

B1^q If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi \Rightarrow \varphi'$, then from $\Delta \hookrightarrow B^r\varphi$ infer $\Delta \hookrightarrow B^r\varphi'$.

B2^q $\Delta \hookrightarrow B^0(true)$.

B3^q From $\Delta \hookrightarrow B^{r_1}\varphi$ and $\Delta \hookrightarrow B^{r_2}\varphi'$ infer $\Delta \hookrightarrow B^{r_3}(\varphi \wedge \varphi')$, where $r_3 = \min(r_1 + r_2, 1)$.

Let $\mathbf{P}^{B,q}_{\Lambda}$ consist of B1^q, B2^q, B3^q,F1$^+$, F3$^+$, EQ, REN, II, and INC. The key reason that a sequence of probability distributions was used even for the quantitative semantics in [1] is that it made it possible to establish an important connection between the qualitative and quantitative semantics. The same argument as that used in [1] for $\mathbf{P}^{+,q}_{\Lambda}$ can be used to show the following result:

Theorem 6. *If* $\mathbf{P}^{B}_{\Lambda} \vdash \Delta \hookrightarrow B\psi$*, then for all* $r \in [0,1]$*, there exists a quantitative instantiation* Δ' *of* Δ *such that* $\mathbf{P}^{B,q}_{\Lambda} \vdash \Delta' \hookrightarrow B^q\psi$*. Moreover,* Δ' *can be found in polynomial time, given the derivation of* $\Delta \hookrightarrow \varphi \rightarrow \psi$*.*

Although we have stated all the axioms above in terms of constants r, r_1, r_2, and r_3, we can easily modify them so as to state them in terms of functions $\epsilon(\eta, tb, \vec{n})$; we leave details to the reader.

This result is quite relevant for us. It says that if we an prove a statement in qualitative conditional logic, then for all desired probability levels $1 - q$, in all models where the assumptions hold with sufficiently high probability, the conclusion will hold with probability at least $1 - q$.

Acknowledgements We thank Riccardo Pucella, Bob Harper and our anonymous reviewers for their comments and suggestions. The authors are supported in part by the AFOSR MURI on Science of Cybersecurity (FA9550-12-1-0040), and the NSF TRUST Center (CCF 0424422). Halpern's work is supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by AFOSR grants FA9550-08-1-0438, FA9550-09-1-0266, and FA9550-12-1-0040, and by ARO grant W911NF-09-1-0281.

References

[1] J. Y. Halpern, "From qualitative to quantitative proofs of security properties using first-order conditional logic," in *Proc. Twenty-Third National Conference on Artificial Intelligence (AAAI '05)*, 2008, pp. 454–459.

[2] M. Goldszmidt, P. Morris, and J. Pearl, "A maximum entropy approach to nonmonotonic reasoning," *IEEE Transactions of Pattern Analysis and Machine Intelligence*, vol. 15, no. 3, pp. 220–232, 1993.

[3] A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turuani, "Probabilistic polynomial-time semantics for a protocol security logic." in *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP '05)*, ser. Lecture Notes in Computer Science. Springer-Verlag, 2005, pp. 16–29.

[4] A. Datta, A. Derek, J. C. Mitchell, and B. Warinschi, "Computationally sound compositional logic for key exchange protocols," in *Proceedings of 19th IEEE Computer Security Foundations Workshop*. IEEE, 2006, pp. 321–334.

[5] C. Fournet, M. Kohlweiss, and P.-Y. Strub, "Modular code-based cryptographic verification," in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 341–350.

[6] W. Diffie, P. C. V. Oorschot, and M. J. Wiener, "Authentication and authenticated key exchanges," *Designs, Codes and Cryptography*, vol. 2, pp. 107–125, 1992.

[7] B. Blanchet, "A computationally sound mechanized prover for security protocols." in *IEEE Symposium on Security and Privacy*, 2006, pp. 140–154.

[8] G. Barthe, B. Grégoire, and S. Zanella Béguelin, "Formal certification of code-based cryptographic proofs," *SIGPLAN Not.*, vol. 44, no. 1, pp. 90–101, Jan. 2009.

[9] G. Barthe, B. Grégoire, S. Heraud, and S. Béguelin, "Computer-aided security proofs for the working cryptographer," in *Advances in Cryptology (CRYPTO 2011)*, ser. Lecture Notes in Computer Science, P. Rogaway, Ed. Springer Berlin Heidelberg, 2011, vol. 6841, pp. 71–90.

[10] A. Roy, A. Datta, and J. C. Mitchell, "Formal proofs of cryptographic security of Diffie-Hellman-based protocols," in *TGC '07: Proceedings of the 3rd Conference on Trustworthy Global Computing*, 2007, pp. 312–329.

[11] S. Goldwasser, S. Micali, and R. L. Rivest, "A digital signature scheme secure against adaptive chosen-message attacks," *SIAM J. Comput.*, vol. 17, no. 2, pp. 281–308, Apr. 1988. [Online]. Available: http://dx.doi.org/10.1137/0217017

[12] M. Bellare and P. Rogaway, "Entity authentication and key distribution," in *Advances in Cryprtology - Crypto '93 Proceedings*. Springer-Verlag, 1994, pp. 232–249.

[13] G. Barthe, J. M. Crespo, B. Grégoire, C. Kunz, Y. Lakhnech, B. Schmidt, and S. Zanella-Béguelin, "Fully automated analysis of padding-based encryption in the computational model," in *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, ser. CCS '13. New York, NY, USA: ACM, 2013, pp. 1247–1260.

[14] G. Bana and H. Comon-Lundh, "Towards unconditional soundness: Computationally complete symbolic attacker," in *Proceedings of the 1st International Conference on Principles of Security and Trust (POST'12)*, ser. Lecture Notes in Computer Science, P. Degano and J. D. Guttman, Eds., vol. 7215. Tallinn, Estonia: Springer, Mar. 2012, pp. 189–208.

[15] D. Dolev and A. Yao, "On the security of public-key protocols," *IEEE Transactions on Information Theory*, vol. 2, no. 29, pp. 198–208, 1983.

[16] M. Abadi and P. Rogaway, "Reconciling two views of cryptography (the computational soundness of formal encryption)." *Journal of Cryptology*, vol. 15, no. 2, pp. 103–127, 2002.

[17] M. Backes, B. Pfitzmann, and M. Waidner, "A universally composable cryptographic library," Cryptology ePrint Archive, Report 2003/015, 2003.

[18] D. Micciancio and B. Warinschi, "Soundness of formal encryption in the presence of active adversaries," in *Proceedings of TCC 2004*, ser. Lecture Notes in Computer Science, vol. 2951. Springer-Verlag, 2004, pp. 133–151.

[19] M. Backes, D. Hofheinz, and D. Unruh, "Cosp: a general framework for computational soundness proofs," in *Proceedings of the 16th ACM Conference on Computer and Communications Security*, ser. CCS '09. New York, NY, USA: ACM, 2009, pp. 66–78.

[20] H. Comon-Lundh and V. Cortier, "Computational soundness of observational equivalence," in *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS'08)*. Alexandria, Virginia, USA: ACM Press, Oct. 2008, pp. 109–118.

[21] V. Cortier and B. Warinschi, "A composable computational soundness notion," in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 63–74.

[22] B. Blanchet, "An Efficient Cryptographic Protocol Verifier Based on Prolog Rules," in *14th IEEE Computer Security Foundations Workshop (CSFW-14)*. Cape Breton, Nova Scotia, Canada: IEEE Computer Society, Jun. 2001, pp. 82–96.

[23] A. Armando, D. A. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron, "The AVISPA tool for the automated validation of internet security protocols and applications." in *Computer Aided Verification, 17th International Conference, CAV 2005, Proceedings*, ser. Lecture Notes in Computer Science, vol. 3576. Springer, 2005, pp. 281–285.

[24] J. Y. Halpern, *Reasoning About Uncertainty*. Cambridge, Mass.: MIT Press, 2003.

[25] M. Bellare and P. Rogaway, "Introduction to modern cryptography," in *Unpublished Lecture Notes*, http://www-cse.ucsd.edu/users/mihir/cse207/classnotes.html.

[26] N. Friedman, J. Y. Halpern, and D. Koller, "First-order conditional logic for default reasoning revisited," *ACM Trans. on Computational Logic*, vol. 1, no. 2, pp. 175–207, 2000.

[27] S. Kraus, D. Lehmann, and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics," *Artificial Intelligence*, vol. 44, pp. 167–207, 1990.